arXiv:1605.00604v1 [cs.SY] 2 May 2016

Formal Verification of Obstacle Avoidance and Navigation of Ground Robots

Stefan Mitsch*, Khalil Ghorbal†, David Vogelbacher‡, André Platzer§

Abstract

The safety of mobile robots in dynamic environments is predicated on making sure that they do not collide with obstacles. In support of such safety arguments, we analyze and formally verify a series of increasingly powerful safety properties of controllers for avoiding both stationary and moving obstacles: (i) *static safety*, which ensures that no collisions can happen with stationary obstacles, (ii) *passive safety*, which ensures that no collisions can happen with stationary or moving obstacles while the robot moves, (iii) the stronger *passive friendly safety* in which the robot further maintains sufficient maneuvering distance for obstacles to avoid collision as well, and (iv) *passive orientation safety*, which allows for imperfect sensor coverage of the robot, i. e., the robot is aware that not everything in its environment will be visible. We complement these provably correct safety properties with *liveness* properties: we prove that provably safe motion is flexible enough to let the robot still navigate waypoints and pass intersections. We use *hybrid system* models and theorem proving techniques that describe and formally verify the robot's discrete control decisions along with its continuous, physical motion. Moreover, we formally prove that safety can still be guaranteed despite sensor uncertainty and actuator perturbation, and when control choices for more aggressive maneuvers are introduced. Our verification results are generic in the sense that they are not limited to the particular choices of one specific control algorithm but identify conditions that make them simultaneously apply to a broad class of control algorithms.

Keywords: provable correctness, obstacle avoidance, navigation, ground robot, hybrid systems

1 Introduction

Autonomous ground robots are increasingly popular as consumer products, ranging from today's autonomous household appliances [7] to the driverless cars of the future being tested on public roads[1]. With the robots leaving the tight confounds of a lab or a locked-off industrial production site, robots face an increased need for ensuring safety not only for the sake of the consumer, but also the manufacturer. One of the most important and challenging safety considerations is to ensure motion safety and that the mobile robot does not collide with any obstacles [3, 37, 39].

In this article, we provide formal proofs as rigorous evidence for the safety of a broad class of obstacle avoidance control algorithms of a robot. In order to take the vagaries of the physical environment into account, these guarantees are for hybrid system models that include discrete control decisions, reaction

*Computer Science Department, Carnegie Mellon University, Pittsburgh, USA,
and Department of Cooperative Information Systems, Johannes Kepler University, Linz, Austria smitsch@cs.cmu.edu
†INRIA, Rennes, France kghorbal@cs.cmu.edu
‡Karlsruhe Institute of Technology, Karlsruhe, Germany uaghg@student.kit.edu
§Computer Science Department, Carnegie Mellon University, Pittsburgh, USA aplatzer@cs.cmu.edu

[1] `http://www.nytimes.com/2010/10/10/science/10google.html?_r=0`

delays, differential equations for the robot's physical motion, bounded sensor uncertainty, and bounded actuator perturbation.

One of the subtle conceptual difficulties, however, is what safety even means for an autonomous robot. We would want it to be collision-free, but that usually requires other vehicles to be sensible, e. g., not actively try to run into the robot when it is just stopped in a corner. One way of doing that is to assume stringent constraints on the behavior of obstacles [3, 14]. In this article, we want to refrain from doing so and allow arbitrary obstacles with an arbitrary continuous motion respecting a known upper bound on their velocity. Then our robot is safe, intuitively, if no collision can ever happen *where the robot is to blame*. For static obstacles, the situation is easy, because the robot is to blame for every collision that happens, so our safety property and its proof show that the robot will never collide with any static obstacle. For dynamic obstacles, safety is subtle, because other moving agents might actively try to ruin safety and cause collisions even if our robot did all it could to prevent them. The first notion we consider is *static safety*, which requires that the robot does not collide with any stationary obstacle. As second notion, we analyze *passive safety* [17], which requires that the robot does not actively collide, i. e., collisions only happen when a moving obstacle ran into the robot while the robot was stopped. Our proofs guarantee passive safety with minimal assumptions about obstacles. The difficulty with passive safety is that it still allows the robot to go kamikaze and stop in unsafe places, creating unavoidable collision situations in which an obstacle has no control choices left that would prevent a collision. The third notion we consider is *passive friendly safety* [17], which aims for more careful robot decisions that respect the features of moving obstacles (e. g., their braking capabilities). A passive friendly robot not only ensures that it is itself able to stop before a collision occurs, but it also maintains sufficient maneuvering room for obstacles to avoid a collision as well. Finally, we introduce *passive orientation safety*, which restricts the responsibility of the robot to avoid collisions to only parts of the robots surroundings (e. g., the robot is responsible for collisions with obstacles to the its front and sides, but obstacles are responsible when hitting the robot from behind).

Motion safety and obstacle avoidance lead to interesting cognitive robotics questions: how much does the robot have to know about the goals and constraints of other vehicles so as not to be considered to blame? In this article, we successively construct models and proofs that increase the level of assumed knowledge and explicitly expressed uncertainty. We start with (i) static safety for the case of static obstacles. Then we consider (ii) passive safety, which assumes a known upper bound on the velocity of obstacles. Then we extend to (iii) passive friendly safety for a known lower bound of an obstacle's braking power and upper bound on its reaction time to initiate collision avoidance attempts. We introduce (iv) a new notion of *passive orientation safety*, which accounts for limited sensor coverage and robot orientation. Finally, we introduce (v) notions of *progress* used in liveness properties to show under which circumstances a safe robot can also reach a goal position.

Note that all our models use symbolic bounds so they hold *for all* choices of the bounds. As a result, we can account for uncertainty in several places (e. g., by instantiating upper bounds on acceleration or time with values including uncertainty). We additionally show how further uncertainty that cannot be attributed to such bounds (in particular location uncertainty, velocity uncertainty, and actuator uncertainty) can be modeled and verified *explicitly*.

The class of control algorithms we consider is inspired by the well-known *dynamic window* algorithm [9], but is equally significant for most other control algorithms when combining our results of provable safety with verified runtime validation [20]. Unlike related work on obstacle avoidance (e. g., [1, 22, 36–38]), we use *hybrid system* models and verification techniques that describe and verify the robot's discrete control choices *along with its continuous, physical motion*. In summary, our contributions are (i) hybrid system models of navigation and obstacle avoidance control algorithms of robots, (ii) safety proofs that they

guarantee static safety, passive safety, passive friendly safety, and passive orientation safety in the presence of stationary and moving obstacles despite sensor uncertainty and actuator perturbation, and (iii) liveness proofs that the safety measures are flexible enough to allow the robot to reach a goal position and pass intersections. The models and proofs of this article are bundled with the KeYmaera theorem prover [34].[2]

The article is organized as follows. In the next section, we discuss related work on navigation and obstacle avoidance of robots that focuses on verification. Section 3 recalls differential dynamic logic that we use as a modeling formalism for the hybrid system dynamics of a robot and the safety and liveness constraints, while Section 4 recalls the dynamic window approach [9] for obstacle avoidance. Section 5 introduces the dynamic model of robots and obstacles that we are going to use throughout this article. In Section 6, we introduce models of obstacle avoidance control and physical motion, and prove that they guarantee static safety, passive safety, passive friendly safety, and passive orientation safety with stationary as well as moving obstacles. We then model uncertainty explicitly and prove that motion is still safe. Section 8 introduces notions of progress and proves liveness of robot navigation. Section 9 illustrates the symbolic findings of the proofs with examples of robots and environments. Section 10 describes how the proofs of this article can be transformed into monitors that provable check safety at runtime, i. e., on a real robot. Section 11 concludes the article.

2 Related Work

Isabelle has recently been used to formally verify that a C program implements the specification of the dynamic window algorithm [37]. This is interesting, but the algorithm itself and its impact on motion of the robot was considered in an informal pen-and-paper argument only. We, instead, formally verify the correctness of the dynamic window algorithm using a hybrid systems verification technique. Our contributions, thus, complement the work in [37] in a twofold manner: First, we create physical models of the control and the motion dynamics of the robot and formally verify correctness of the dynamic window algorithm control for the combined hybrid systems dynamics. Second, we model stationary as well as moving obstacles and prove multiple safety properties. These complementary results together present a strong safety argument from concept (this article) to implementation [37].

PASSAVOID [3] is a navigation scheme, which avoids *braking inevitable collision states* (i. e., states that regardless of the robot's trajectory lead to a collision) to achieve safety in the presence of moving obstacles. Since PASSAVOID is designed to operate in completely unknown environments, it ensures that the robot is at rest when a collision occurs (*passive safety*). The motion dynamics of the robot have only been considered in simulation. We prove the stronger *passive friendly safety* using a hybrid verification technique (i. e., algorithm and motion dynamics), which ensures that the robot does not create unavoidable collision situations by stopping in unsafe places.

For the purpose of guaranteeing infinite horizon safety, velocity obstacle sets [39] assume unpredictable behavior for obstacles with known forward speed and maximum turn rate (i. e., Dubin's cars). The authors focus only on the obstacle behavior; the robot's motion is explicitly excluded from their work. We complement their work and show that a robot, which has a known upper bound on its reaction time and considers discs as velocity obstacle sets (i. e., known forward speed and unknown turn rate, as allowed by [39]), moves safely.

Hybrid system models of driver support systems in cars [14, 19] have been verified with a model of the continuous dynamics of cars. That points out interesting safety conditions for vehicles on straight lines, but not in the general motion in the two-dimensions plane that we consider in this work.

[2] `http://symbolaris.com/info/KeYmaera.html`

Safety of aircraft collision avoidance maneuvers in the two-dimensional plane was verified for constant translational velocity and a rotational velocity that stays constant during the maneuver [15, 33]. Our models include acceleration for both translational and rotational velocity and are generalized to address uncertainty.

LTLMoP contains an approach [35] to study high-level behavior for map exploration when the environment is continuously updated. The approach synthesizes and re-synthesizes plans, expressed in linear temporal logic, of a hybrid controller, when new map information is discovered. This work focuses on preserving the state and task completion history, and thus on guaranteeing that the robot will follow a high-level behavior (e. g., visit all rooms) even when the controller is re-synthesized, not on safe obstacle avoidance.

Pan et al. [22] propose a method to smooth the trajectories produced by sampling-based planners in a collision-free manner. Our article proves that such trajectories are indeed safe when considering the control choices of a robot and its continuous dynamics.

LQG-MP [38] is a motion planning approach that takes into account the sensors, controllers, and motion dynamics of a robot while working with uncertain information about the environment. The approach assesses randomly generated paths by the approximated probability of a collision with an obstacle. One goal is to select paths with low collision probability; however, guaranteeing collision-free motion is not their focus, since a collision-free path may not even have been generated.

Althoff et al. [1] use a probabilistic approach to rank trajectories according to their collision probability. To further refine such a ranking, a collision cost metric is proposed, which derives the cost of a potential collision by considering the relative speeds and masses of the collision objects.

Seward et al. [36] try to avoid potentially hazardous situations by using Partially Observable Markov Decision Processes. Their focus, however, is on a user-definable trade-off between safety and progress towards a goal. Hence, safety is not guaranteed under all circumstances.

In summary, this article addresses safety of robot obstacle avoidance in the following manner.

- Unlike [3, 35–37, 39], we study combined models of the *hybrid system dynamics* in terms of discrete control and differential equations for continuous physical motion of the robot as well as the obstacles, not only discrete control alone or only the behavior of obstacles.
- Unlike [14, 15, 19, 33, 37], we verify safety in the two-dimensional plane not one-dimensional space *and* do not assume constant translational and rotational velocity, but include accelerations for both, as needed for ground vehicles.
- Unlike [1, 22, 35, 36, 38], we produce *formal, deductive proofs* in a formal verification tool. Note, that in [37] the correctness of the implementation w.r.t. the algorithm's specification was formally verified, but not the motion of the robot.
- Unlike [22, 35, 37], we verify safety even in the presence of moving obstacles.
- Unlike [3, 22, 35, 37], we verify *passive friendly safety*, which is important because passive (non-friendly) safe robots may cause unavoidable collisions by stopping in unsafe places so that obstacles will collide with them. We also verify passive orientation safety, which is important for situations with limited sensor ranges.
- Unlike [3, 14, 33, 35], we consider sensor and actuator uncertainty in our verification results.
- Unlike [1, 36, 38], we do not minimize or probabilistically minimize collisions, but prove that collisions can never occur (as long as the robot fits to the model).

3 Preliminaries: Differential Dynamic Logic

A robot and the moving obstacles in its environment form a hybrid system: they make discrete control choices (e. g., compute the actuator set values for acceleration, braking, or steering), which in turn influence their actual physical behavior (e. g., slow down to a stop, move along a curve). In a test-driven approach, a simulator or field tests provide insight into the expected physical effects of the control code. In formal verification, by analogy, hybrid systems have been considered as joint models for both discrete and continuous behavior, since verification of either component alone does not capture the full behavior of a robot and its environment. In this section, we first give an overview of the relationship between testing, simulation, and formal verification, before we introduce the syntax and semantics of the specification language that we use for formal verification.

3.1 Testing, Simulation, and Formal Verification

Testing, simulation, and formal verification complement each other. Testing helps to make a system robust under real-world conditions, whereas simulation lets us execute a large number of tests in an inexpensive manner (at the expense of realistic conditions). Both, however, show correctness for the tested scenarios only: testing and simulation discover the presence of bugs, but cannot show their absence. Formal verification, in contrast, provides precise and undeniable guarantees for all possible executions of the modeled behavior, assuming that the model adequately captures reality. Formal verification either discovers bugs if present, or shows the absence of bugs in the model, but cannot show whether or not the model is realistic. In Section 10, we will see how we can use runtime monitoring to bridge both worlds.

Testing, simulation, and formal verification all base on similar ingredients, but apply different levels of rigor:

Software Testing and simulation run a specific control algorithm with specific parameters (e. g., run some specific version of an obstacle avoidance algorithm with maximum velocity $V = 2\,\text{m/s}$); in formal verification, we can specify symbolic parameters and nondeterministic inputs and effects and, thereby, capture entire families of control algorithms and many scenarios at once (e. g., verify all velocities $0 \leq v \leq V$ for any maximum velocity V at once).

Hardware and physics Testing runs a real robot in a real environment. Both simulation and formal verification, in contrast, base on a model of the hardware and physics to provide sensor values and compute how software decisions result in real-world effects.

Requirements Testing and simulation can work with informal or semi-formal requirements (e. g., a robot should not collide with obstacles—allows slack: could mean that collisions are acceptable when other outcomes are even less desirable, whatever that means), whereas formal verification bases on mathematically precise formal requirements expressed as a logical formula (allows no slack in interpretation: unambiguously distinguishes between correct behavior and faults).

Process In testing and simulation, requirements are formulated as test conditions and expected test outcomes. A test procedure then runs the robot several times under the test conditions and compares the actual output with the expected outcome (e. g., run the robot in different spaces, with different obstacles, various software parameters, and different sensor configurations to see whether or not any of the runs fail to avoid obstacles). The test protocol serves as correctness evidence. In formal verification, the requirements are formulated as a logical formula. A theorem prover then creates a mathematical

proof showing that all possible executions—possibly infinitely many—of the model are correct (safety proof), or showing that the model has a way to achieve a goal (liveness proof). The mathematical proof is a correctness certificate.

The remainder of this section focuses on the syntax and semantics that we use for formal verification: it introduces *hybrid programs*, which is a program notation for describing hybrid systems, and *(quantified) differential dynamic logic* for specifying correctness conditions about these programs. Using hybrid programs, we can specify how the robot and the obstacles in the environment make decisions and move physically. With differential dynamic logic we can specify formally which behavior of a hybrid program is considered correct.

3.2 Differential Dynamic Logic

In order to verify safe obstacle avoidance, we use *differential dynamic logic* $\mathsf{d}\mathcal{L}$ [23, 25, 29], which has a notation for hybrid systems as *hybrid programs*. We use hybrid programs for modeling a robot that follows the dynamic window algorithm as well as for modeling the behavior of moving obstacles. $\mathsf{d}\mathcal{L}$ allows us to make statements that we want to be true for all runs of the program (safety) or for at least one run (liveness). Both constructs are necessary to verify safety: for all possible control choices and entailed physical motion, our robot must be able to stop, while at the same time there must be at least one possible execution in which the obstacle is able to stop without collision as well.

One of the many challenges of developing robots is that we do not know the behavior of the environment exactly. For example, a moving obstacle may or may not slow down when our robot approaches it. In addition to programming constructs familiar from other languages (e. g., assignments and conditional statements), hybrid programs, therefore, provide nondeterministic operators that allow us to describe such unknown behavior of the environment concisely. These nondeterministic operators are also useful to describe parts of the behavior of our own robot (e. g., we may not be interested in the exact value delivered by a position sensor, but only that it is within some error range), which then corresponds to verifying an entire family of controllers. Using nondeterminism to model our own robot has the benefit that later optimization (e. g., mount a better sensor or implement a faster algorithm) does not necessarily require re-verification.

Table 1 summarizes the syntax of hybrid programs together with an informal semantics. Below, we briefly describe each operator with an example. Sequential composition $\alpha;\beta$ says that β starts after α finishes (e. g., first let the robot choose acceleration, then steering angle). The nondeterministic choice $\alpha \cup \beta$ follows either α or β (e. g., the obstacle may or may not slow down). The nondeterministic repetition operator α^* repeats α zero or more times (e. g., the robot may encounter obstacles over and over again, but we do not know exactly how often). Assignment $x := \theta$ instantaneously assigns the value of the term θ to the variable x (e. g., let the robot choose maximum braking), while $x := *$ assigns an arbitrary value to x (e. g., an obstacle may choose any acceleration, we do not know which value exactly). $x' = \theta \,\&\, F$ describes a continuous evolution of x within the evolution domain F (e. g., let the velocity of the robot decrease according to the applied brakes, but not become negative—hitting the brakes won't make the robot drive backwards). The test $?F$ checks that a particular condition F holds, and aborts if it does not (e. g., test whether or not the distance to an obstacle is large enough to continue with accelerating). A typical pattern that involves assignment and tests is to limit the assignment of arbitrary values to known bounds (e. g., limit an arbitrarily chosen acceleration to the physical limits of the robot, as in $x := *;\ ?x \leq A$, which says x is any value less or equal A).

The set of $\mathsf{d}\mathcal{L}$ formulas is generated by the following EBNF grammar (where $\sim \in \{<, \leq, =, \geq, >\}$

Table 1: Hybrid program representations of hybrid systems.

Statement	Effect
$\alpha;\ \beta$	sequential composition, first run α, then β
$\alpha \cup \beta$	nondeterministic choice, following either α or β
α^*	nondeterministic repetition, repeats α $n \geq 0$ times
$x := \theta$	assign value of term θ to variable x (discrete jump)
$x := *$	assign arbitrary real number to variable x
$\big(x_1' = \theta_1, \ldots,$	evolve x_i along differential equation system $x_i' = \theta_i$
$\quad x_n' = \theta_n \,\&\, F\big)$	restricted to maximum evolution domain F

and θ_1, θ_2 are arithmetic expressions in $+, -, \cdot, /$ over the reals):

$$\phi ::= \theta_1 \sim \theta_2 \mid \neg\phi \mid \phi \wedge \psi \mid \phi \vee \psi \mid \phi \rightarrow \psi \mid \forall x \phi \mid [\alpha]\phi \mid \langle\alpha\rangle\phi$$

Further operations, such as Euclidian norm $\|\theta\|$ and infinity norm $\|\theta\|_\infty$ of a vector θ, are definable.

To specify the desired correctness properties of hybrid programs, dL formulas of the form $F \rightarrow [\alpha]G$ mean that all executions of the hybrid program α, which start at a state in which formula F is true, lead to states in which formula G is true. Dually, formula $F \rightarrow \langle\alpha\rangle G$ expresses that there is a state reachable by the hybrid program α that satisfies formula G.

Differential dynamic logic comes with a verification technique to prove those correctness properties. We did all our proofs in the verification tool KeYmaera [34], which implements this verification technique [23, 25, 30]. KeYmaera supports hybrid systems with nonlinear discrete jumps, nonlinear differential equations, differential-algebraic equations, differential inequalities, and systems with nondeterministic discrete or continuous input. This makes KeYmaera more readily applicable to robotic verification than other hybrid system verification tools, such as SpaceEx [10], which focuses on piecewise linear systems. KeYmaera implements automatic proof strategies that decompose hybrid systems symbolically [25]. This compositional verification principle helps scaling up verification, because KeYmaera verifies a big system by verifying properties of subsystems. Strong theoretical properties, including relative completeness results, have been shown about dL [23, 29].

3.3 Quantified Differential Dynamic Logic

Often, we want to analyze how the robot interacts with many other agents (e. g., avoid collision with each one of many obstacles) or represent that it consists of or uses multiple instances of some device (e. g., several location sensors). In order to prevent duplicating variables for each of the objects, which is undesirable even for a very small, known number of objects, we need a way of referring to countably many objects concisely. In this article, we discuss two ways of referring to countably many objects with differential dynamic logic:

- we can implicitly refer to many obstacles by choosing one nondeterministically (see nondeterministic assignment above, used in Section 6.1–7.4), or
- we can explicitly refer to each obstacle individually by using quantification over objects of a sort (e. g., each object of the sort obstacle, used in Section 7.5).

Quantified differential dynamic logic QdL [26, 28] is an extension of differential dynamic logic suited for verifying distributed hybrid systems by quantifying over sorts. The notion of hybrid programs is extended

to quantified hybrid programs. Instead of using a single state variable x to describe an attribute of an object, we can use a function term $x : O \to \mathbb{R}$ in QdL to denote that object i has x-value $x(i)$, for each i of sort O. We use a non-rigid function term if we want to change its value (e. g., the position of an obstacle); otherwise it is rigid (e. g., the maximum velocity of an obstacle). A sort describes hereby a class of objects (e. g., a sort could be the class of all obstacles or the class of all moving obstacles). Pure differential dynamic logic dL uses the sort $\mathbb{R}$. An alternative way of expressing sorts would be to use a unary predicate symbol that is true if and only if the object is of the desired sort. In QdL formulas we can use quantifiers to make statements about the instances of a sort S with $\forall i \in S$ and $\exists i \in S$, similar to what is possible in dL for the special sort $\mathbb{R}$.

QdL allows us to explicitly track properties of all obstacles simultaneously. Quantified hybrid programs allow the evolution of properties expressed as non-rigid functions for all objects of the same sort simultaneously (e. g., all obstacles move simultaneously).

Table 2 lists statements that can be used in quantified hybrid programs in addition to those of hybrid programs [26, 28].

Table 2: Statements of quantified hybrid programs [26, 28].

Statement	Effect
$\forall i \in C\ x(i) := \theta$	Assigns the value of term θ to x of all objects of sort C. x is a non-rigid function.
$\forall i \in C\ \big(x_j(i)' = \theta_j(i)\ \&\ F\big)$	evolve all $x(i)$ along differential equations $x_j(i)' = \theta_j(i), \ldots$ restricted to evolution domain F

4 Preliminaries: Obstacle Avoidance with the Dynamic Window Approach

The robotics community has come up with an impressive variety of robot designs, which differ not only in their tool equipment, but also (and more importantly for the discussion in this article) in their kinematic capabilities. We focus on wheel-based vehicles. In order to make our models applicable to a large variety of robots, we use only limited control options (e. g., do not move sideways to avoid collisions since Ackermann drive could not follow such evasion maneuvers). We consider robots that drive forward (non-negative translational velocity) in sequences of arcs in two-dimensional space.[3] Such trajectories can be realized by robots with single-wheel drive, differential drive, Ackermann drive, synchro drive, or omni drive [4]. In a nutshell, in order to stay on the safe side, our models conservatively underestimate the capabilities of our robot while conservatively overestimating the dynamic capabilities of obstacles.

Many different navigation and obstacle avoidance algorithms have been proposed for such robots, e. g. *dynamic window* [9], *potential fields* [12], or *velocity obstacles* [8]. For an introduction to various navigation approaches for mobile robots, see [2, 6]. In this article, our focus is on the dynamic window algorithm [9], which is derived from the motion dynamics of the robot and thus discusses all aspects of a hybrid system (models of discrete and continuous dynamics). Other control algorithms including even path planners based on RRT [13] or A* [11] are compatible with our results when their control decisions are checked with a runtime verification approach [20] against the safety conditions we identify for the motion here.

[3] If the radius of such a circle is infinite, the robot drives (forward) on a straight line.

Table 3: Parameters, state variables of robot and obstacle

Notation	Coordinates	Description
p_r	(p_r^x, p_r^y)	Position of the robot
v_r		Translational velocity
a_r		Translational acceleration, s.t. $-b \leq a_r \leq A$
ω_r		Rotational velocity, s.t. $\omega_r r_c = v_r$
d_r	(d_r^x, d_r^y)	Orientation of the robot, s.t. $\|d_r\| = 1$
p_c	(p_c^x, p_c^y)	Curve center, s.t. $d_r = (p_r - p_c)^\perp$
r_c		Curve radius, s.t. $r_c = \|p_r - p_c\|$
p_o	(p_o^x, p_o^y)	Position of the obstacle
v_o	(v_o^x, v_o^y)	Translational velocity, including orientation, s.t. $\|v_o\| \leq V$
A		Maximum acceleration $A \geq 0$
b		Minimum braking $b > 0$
ε		Maximum control loop reaction delay $\varepsilon > 0$
V		Maximum obstacle velocity $V \geq 0$
Ω		Maximum rotational velocity $\Omega \geq 0$

The dynamic window algorithm is an obstacle avoidance approach for mobile robots equipped with synchro drive [9] but can be used for other drives too [5]. It uses circular trajectories that are uniquely determined by a translational velocity v together with a rotational velocity ω. The algorithm is roughly organized into two steps: (i) The range of all possible pairs of translational and rotational velocities is reduced to admissible ones that result in safe trajectories (i. e., avoid collisions since those trajectories allow the robot to stop before it reaches the nearest obstacle). The admissible pairs are further restricted to those that can be realized by the robot within a short time frame (the so-called dynamic window). If the set of admissible and realizable velocities is empty, the algorithm stays on the previous safe curve (such curve exists unless the robot started in an unsafe state). (ii) Progress towards the goal is optimized by maximizing a goal function among the set of all admissible controls. For safety verification, we can omit the second step and verify the stronger property that *all* choices that are fed into the optimization are safe, since even if none is identified, the previous safe curve can be continued.

5 Robot and Obstacle Motion Model

This section introduces the robot and obstacle motion models that we are going to use throughout the article. Table 3 summarizes the model variables and parameters of both the robot and the obstacle for easy reference. In the following subsections, we illustrate their meaning in detail.

5.1 Robot State and Motion

The dynamic window algorithm safely abstracts the robot's shape to a single point, since other shapes reduce to adjusting the (virtual) shapes of the obstacles (cf. [18] for an approach to attribute robot shape to obstacles). We also use this abstraction to reduce the verification complexity. Fig. 1 illustrates how we model the position, orientation, and trajectory of a robot.

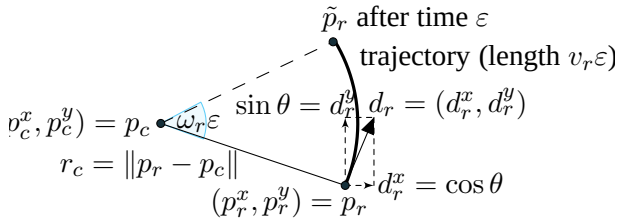

Figure 1: State illustration of a robot on a two-dimensional plane. The robot has position $p_r = (p_r^x, p_r^y)$, orientation $d_r = (d_r^x, d_r^y)$, and drives on circular arcs (thick arc) of radius r_c with translational velocity v_r, rotational velocity ω_r and thus angle $\omega_r\varepsilon$ around curve center points $p_c = (p_c^x, p_c^y)$. In time ε the robot will reach a new position $\tilde{p}_r$, which is $v_r\varepsilon$ away from the initial position p_r when measured along the robot's trajectory arc.

The robot has state variables describing its current position $p_r = (p_r^x, p_r^y)$, translational velocity $v_r \geq 0$, translational acceleration a_r, an orientation vector[4] $d_r = (\cos\theta, \sin\theta)$, and angular velocity[5] $\theta' = \omega_r$. The translational and rotational velocities are linked w.r.t. the rigid body planar motion by the formula $r_c\omega_r = v_r$, where the curve radius $r_c = \|p_r - p_c\|$ is the distance between the robot and the center of its current curve $p_c = (p_c^x, p_c^y)$. Following [24], we use differential axiomatization to encode sine and cosine functions in the dynamics using the extra variables $d_r^x = \cos\theta$ and $d_r^y = \sin\theta$ to avoid undecidable arithmetic. The continuous dynamics for the dynamic window algorithm [9] can be described by the differential equation system of ideal-world dynamics of the planar rigid body motion:

$$p_r' = v_r d_r, v_r' = a_r, d_r' = \omega_r d_r^\perp, (r_c\omega_r)' = a_r$$

where

- $p_r' = v_r d_r$ is vector notation for $p_r^{x\prime} = v_r d_r^x, p_r^{y\prime} = v_r d_r^y$,
- the condition $d_r' = \omega d_r^\perp$ is vector notation for $d_r^{x\prime} = -\omega_r d_r^y, d_r^{y\prime} = \omega_r d_r^x$ where $^\perp$ is the orthogonal complement, and
- the condition $(r_c\omega_r)' = a_r$ encodes the rigid body planar motion $r_c\omega_r = v_r$ that we consider.

The dynamic window algorithm assumes direct instantaneous control of the translational velocity v_r. We, instead, control acceleration a_r and do not perform instant changes of the velocity. Our model is closer to the actual dynamics of a robot, which cannot really change its velocity from 20 to 2 instantly. The realizable velocities follow from the differential equation system according to the controlled acceleration a_r.

Fig. 2a depicts the position and velocity changes of a robot accelerating on a circle around a center point $p_c = (2, 0)$. The robot starts at $p_r = (0, 0)$ as initial position, with $v_r = 2$ as initial translational velocity and $\omega_r = 1$ as initial rotational velocity; Fig. 2d shows the resulting circular trajectory. Fig. 2b and Fig. 2e show the resulting curve when braking (the robot brakes along the curve and comes to a complete stop before completing the circle). If the rotational velocity is constant ($\omega_r' = 0$), the robot drives an Archimedean

[4] As stated earlier, we study unidirectional motion: the robot moves along its direction, that is the vector d_r gives the direction of the velocity vector.

[5] The derivative with respect to time is denoted by prime ($'$).

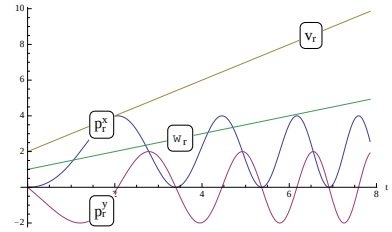

(a) Position (p_r^x, p_r^y), translational velocity v_r and rotational velocity ω_r for positive acceleration on a circle.

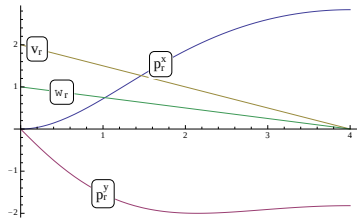

(b) Position (p_r^x, p_r^y), translational velocity v_r and rotational velocity ω_r for braking to a complete stop on a circle.

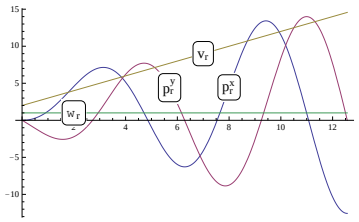

(c) Position (p_r^x, p_r^y), translational velocity v_r and rotational velocity ω_r for translational acceleration on a spiral.

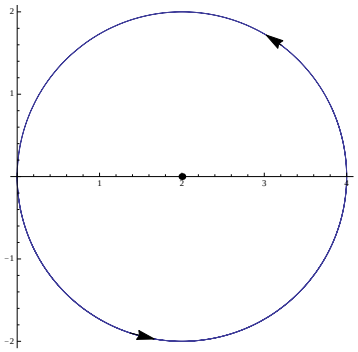

(d) (p_r^x, p_r^y) motion plot for acceleration a.

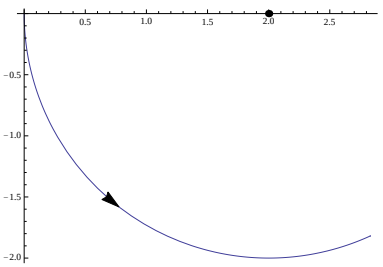

(e) (p_r^x, p_r^y) motion plot for braking b.

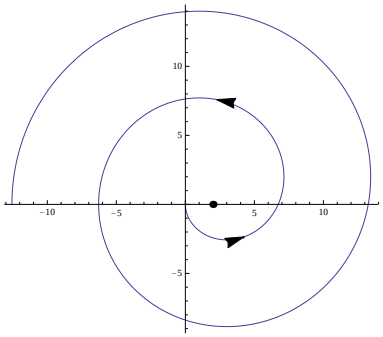

(f) (p_r^x, p_r^y) motion plot for c.

Figure 2: Trajectories of the robot over time (top) or in planar space (bottom).

spiral with the translational and rotational accelerations controlling the spiral's separation distance (a_r/ω_r^2). The corresponding trajectories are shown in Figures 2c and 2f.

We assume bounds for the permissible acceleration a_r in terms of a maximum acceleration $A \geq 0$ and braking power $b > 0$, as well as a bound Ω on the permissible rotational velocity ω_r. We use ε to denote the upper bound for the control loop time interval (e. g., sensor and actuator delays, sampling rate, and computation time). That is, the robot may react as quickly as it wants, but it can take no longer than time ε. Note, that the robot would not be safe without such a time bound, because its control might never run. In our model, all these bounds will be used as symbolic parameters and not concrete numbers. Therefore, our results apply to *all values* of these parameters and can be enlarged to include uncertainty.

5.2 Obstacle State and Motion

An obstacle has (vectorial) state variables describing its current position $p_o = (p_o^x, p_o^y)$ and velocity $v_o = (v_o^x, v_o^y)$. The obstacle model is deliberately very liberal to account for many different obstacle behaviors. The only restriction about the dynamics is that the obstacle moves continuously with bounded velocity $\|v_o\| \leq V$ while the physical system evolves for ε time units. Note, that the dynamic window algorithm considers a special case $V = 0$ (obstacles are stationary). Depending on the relation of V to ε, moving obstacles can make quite a difference, e. g., when fast obstacles meet communication-based virtual sensors as in RoboCup.[6]

[6] `http://www.robocup.org/`

6 Safety Verification of Ground Robot Motion

Table 4: Overview of safety notions, responsibilities of the robot and its assumptions about the obstacle

Safety	**Responsibility of Robot**	**Assumptions about Obstacles**
Static Section 6.1	Positive distance to all stationary obstacles	Obstacles remain stationary and never move
	$\psi_{\text{ss}} \equiv \|p_r - p_o\| > 0$	$\zeta_{\text{ss}} \equiv v_o = 0$
Passive Section 6.2	Positive distance to all obstacles while driving	Known maximum velocity of obstacles
	$\psi_{\text{ps}} \equiv v \neq 0 \rightarrow \|p_r - p_o\| > 0$	$\zeta_{\text{ps}} \equiv \|v_o\| \leq V$
Passive Friendly Section 6.3	Passive safety plus sufficient maneuvering space for obstacles when stopped	Passive safety plus known minimum braking capability and known maximum reaction time of obstacles
	$\psi_{\text{pfs}} \equiv \psi_{\text{ps}} \wedge (v_r = 0 \wedge \|p_r - p_o\| > \frac{V^2}{2b_o} + \tau V \wedge 0 \leq v_o \leq V) \rightarrow \langle \text{obstacle} \rangle (\|p_r - p_o\| > 0 \wedge v_o = 0)$	$\zeta_{\text{pfs}} \equiv \zeta_{\text{ps}} \wedge b_o > 0 \wedge \tau \geq 0$
Passive Orientation Section 6.4	Positive distance to all obstacles while driving, unless an invisible obstacle interfered with the robot while the robot stayed cautiously inside its observable region	Passive safety
	$v_r \neq 0 \rightarrow \big(\|p_r - p_o\| > 0 \vee (\mathit{isVisible} \leq 0 \wedge \|\beta\| < \gamma)\big)$	ζ_{ps}

We want to prove motion safety of a robot that avoids obstacles by dynamic window navigation. Starting from a simplified robot controller, we develop increasingly more realistic models, and discuss different safety notions. *Static safety* describes a vehicle that never collides with stationary obstacles. Under *passive safety* [17], the vehicle is in a safe state if it is able to come to a full stop before making contact with an obstacle (i. e., the vehicle does not itself collide with obstacles, so if a collision occurs at all then while the vehicle was stopped). Passive safety, however, puts the burden of avoiding collisions mainly on other objects. We further want to prove the stronger *passive friendly safety* [17]: we want to guarantee that our robot will come to a full stop safely under all circumstances, but will also leave sufficient maneuvering room for moving obstacles to avoid a collision.[7] Finally, we want to prove passive orientation safety, which accounts for the sensor coverage of the robot and its orientation to reduce the responsibility of the robot in structured spaces, such as on roads with lanes.

Table 4 gives an overview of the safety notions (both formally and informally) and the assumptions made about the robot and the obstacle in the models. We consider all four models and safety properties to show the differences between the assumed knowledge and the safety guarantees that can be made. The

[7] The robot ensures that there is enough room for the obstacle to stop before a collision occurs. If the obstacle decides not to, the obstacle is to blame and our robot is still considered safe.

verification effort and complexity difference is quite instructive. Static safety provides a strong guarantee with a simple safety proof, because only the robot moves. Passive safety can be guaranteed by proving safety of all robot choices, whereas passive friendly safety additionally requires liveness proofs for the obstacle. In the following sections, we discuss models and verification of the dynamic window algorithm in detail.

For the sake of clarity, we will initially make the following simplifying assumptions to get an easier first model:

- in its decisions, the robot will assume it uses maximum braking or maximum acceleration,
- the robot will not be able to reverse its direction, but only drive smooth curves in forward direction, and
- the robot will not keep track of the center of the circle around which its current arc is taking it, but choose steering through picking a curve radius

In Section 7 we will see how to avoid these simplifications.

The subsections here are structured as follows: we first discuss the rationale behind the model (see paragraphs *Modeling*) and give an intuition why the control choices in this model are safe (see paragraphs *Identification of Safe Controls*). Finally, we verify the correctness of the model formally, i. e., we use the model in a correctness theorem and sketch a proof that the control choices indeed guarantee the model to satisfy the static safety condition (see paragraphs *Verification*). Whether or not the model adequately represents reality is a complementary question that we will discuss in Section 10.

6.1 Static Safety with Maximum Acceleration

In environments with only stationary obstacles, static safety ensures that the robot will never collide.

Modeling The prerequisite for obtaining a formal safety result is to first formalize the system model in addition to its desired safety property. We develop a model of the principles in the dynamic window algorithm as a hybrid program, and express static safety as a safety property in $\mathsf{d}\mathcal{L}$.

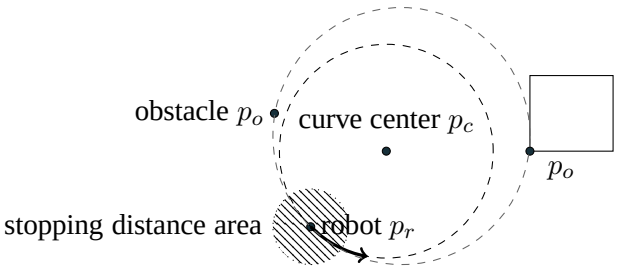

Figure 3: Illustration of static safety: the robot must stop before reaching the closest obstacle on a curve (two of infinitely many curves illustrated). We abstract non-point obstacles to points by considering the perimeter point being closest to the robot. A conservative simplification is to ignore the curves of the robot and use a safety zone of the size of its stopping distance instead (dotted circle).

The dynamic window algorithm uses the distance to the nearest obstacle for every possible curve to determine admissible velocities (e. g., compute distances in a loop and pick the obstacle with the smallest). Our model exploits the power of nondeterminism to model this concisely. It nondeterministically picks *any*

obstacle $p_o := (*, *)$ and tests its safety. Since the choice of the obstacle to consider was nondeterministic and the model is only safe if it is safe for *all* possible ways of selecting *any* obstacle nondeterministically, this includes safety for the nearest obstacle (ties are included) and is thus safe for all possible obstacles. A QdL model with explicit representations of multiple obstacles will be considered in Section 7.5. In the case of non-point obstacles, p_o denotes the obstacle perimeter point that is closest to the robot (this fits naturally to obstacle point sets delivered by radar and Lidar sensors, from which the closest point will be chosen). In each controller run of the robot, the position p_o is updated nondeterministically (again to consider any obstacle including the ones that are now closest). In this process, the robot may or may not find another safe trajectory. If it does, the robot can follow that new safe trajectory w.r.t. any nondeterministically chosen obstacle. If not, the robot can still brake on the previous trajectory, which was previously shown to be safe.

Model 1 Dynamic window with static safety

$$\textit{dw}_{\text{ps}} \equiv (\textit{ctrl}_r; \textit{dyn})^* \tag{1}$$
$$\textit{ctrl}_r \equiv \quad (a_r := -b) \tag{2}$$
$$\cup\ (?v_r = 0;\ a_r := 0;\ \omega_r := 0) \tag{3}$$
$$\cup\ (a_r := A;\ \omega_r := *;\ ?{-\Omega} \leq \omega_r \leq \Omega; \tag{4}$$
$$r_c := *;\ p_o := (*, *);\ ?\textit{curve} \wedge \textit{safe}) \tag{5}$$
$$\textit{curve} \equiv r_c \neq 0 \wedge r_c\omega_r = v_r \tag{6}$$
$$\textit{safe} \equiv \|p_r - p_o\|_\infty > \frac{v_r^2}{2b} + \left(\frac{A}{b} + 1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon v_r\right) \tag{7}$$
$$\textit{dyn} \equiv t := 0;\ \{t' = 1,\ {p_r}' = v_r d_r,\ v_r' = a_r, \tag{8}$$
$$d_r' = \omega_r d_r^\perp,\ \omega_r' = \frac{a_r}{r_c} \tag{9}$$
$$\&\ v_r \geq 0 \wedge t \leq \varepsilon\} \tag{10}$$

Model 1 represents the common controller-plant model: it repeatedly executes the robot control choices followed by dynamics, cf. (1). The continuous dynamics of the robot as presented in Section 5 above are defined in (8)–(10) of Model 1. The rest describes the discrete control. For the sake of clarity we restrict the study to circular trajectories with non-zero radius (that is $r_c \neq 0$ so that the robot is not spinning on the spot), where straight-line trajectories correspond to infinite r_c. The sign of the radius signifies if the robot follows the curve in clockwise ($r_c < 0$) or in counter-clockwise direction ($r_c > 0$). Since $r_c \neq 0$, the condition $(r_c\omega_r)' = a_r$ can be rewritten as $\omega_r' = \frac{a_r}{r_c}$.

The robot is allowed to brake at all times since the assignment that assigns full braking to a_r in (2) has no test. If the robot is stopped, it may choose to stay in its current spot without turning, cf. (3). Finally, if it is safe to accelerate, the robot may choose a new safe curve in its dynamic window: it chooses maximum acceleration, and any rotational velocity in the bounds, cf. (4). This corresponds to testing all possible rotational velocity values at the same time. An implementation in an imperative language would use loops to enumerate all possible values and all obstacles and test each pair (v_r, ω_r) separately w.r.t. every obstacle, storing the admissible pairs in a data structure (as e. g., in [37]).

The curve is determined by the robot following a circular trajectory of radius r_c, starting in initial direction d_r with angular velocity ω_r, cf. (5). The distance to the nearest obstacle on that curve is measured through $p_o := (*, *)$ in (5). The trajectory starts at p_r with translational velocity v_r and rotational velocity

ω_r, as defined by $r_c\omega_r = v_r$ in (6). Together with the orientation d_r of the robot, which is tangential to the curve, this implicitly defines the rotation center p_c; see Fig. 3. We will explicitly represent the rotation center in later models for more aggressive maneuvering; for starters, here, we only need to know how to steer through r_c.

Identification of Safe Controls Based on this shape of the model, its most critical element is the formula *safe* that we identify as the condition that control choices need to satisfy in order to always keep the robot safe. While its ultimate justification will be the safety proof (Theorem 1), this section develops an intuitive explanation why we chose the particular design in (7).

A circular trajectory of radius r_c ensures static safety if it allows the robot to stop before it collides with the nearest obstacle. Consider the extreme case where the radius $r_c = \infty$ is infinitely large and the robot, thus, travels on a straight line. In this case, the distance between the robot's current position p_r and the nearest obstacle p_o must account for the following components: First, the robot needs to be able to brake from its current velocity v_r to a complete stop:

$$\tfrac{v_r^2}{2b} = \int_0^{v_r/b} (v_r - bt)dt \ . \tag{11}$$

Second, it may take up to ε time until the robot can take the next control decision. Thus, we must additionally take into account the distance that the robot may travel w.r.t. the maximum acceleration A and the distance needed for compensating its acceleration of A during that reaction time with braking power b:

$$\left(\tfrac{A}{b}+1\right)\left(\tfrac{A}{2}\varepsilon^2 + \varepsilon v_r\right) = \int_0^{\varepsilon} (v_r + At)dt + \int_0^{A\varepsilon/b} (v_r + A\varepsilon - bt)dt \ . \tag{12}$$

The safety distance chosen for *safe* in (7) of Model 1 is the sum of the distances (11) and (12). The safety proof will have to show that this construction was safe and that it is also safe for all other curved trajectories that the obstacle and robot could be taking instead.

To simplify the proof's arithmetic, we measure the distance between the robot's position p_r and the obstacle's position p_o in the infinity-norm $\|p_r - p_o\|_\infty$, i. e., either $|p_r^x - p_o^x|$ or $|p_r^y - p_o^y|$ must be safe. In the illustrations, this corresponds to replacing the circles representing reachable areas with squares. We, thus, over-approximate the Euclidean norm distance $\|p_r - p_o\|_2 = \sqrt{(p_r^x - p_o^x)^2 + (p_r^y - p_o^y)^2}$ by a factor of at most $\sqrt{2}$.

Verification We verify the safety of the control algorithm modeled as a hybrid program in Model 1, using a formal proof calculus for dL [23, 25, 30]. The robot is safe, if it maintains positive distance $\|p_r - p_o\| > 0$ to all obstacles (see Table 4):

$$\psi_{\text{ss}} \ \equiv \ \|p_r - p_o\| > 0 \ . \tag{13}$$

In order to guarantee ψ_{ss}, the robot must stay at a safe distance, which still allows the robot to brake to a complete stop before hitting any obstacle. The following condition captures this requirement as an invariant φ_{ss} that we prove to hold for all loop executions:

$$\varphi_{\text{ss}} \ \equiv \ \|p_r - p_o\| > \frac{v_r^2}{2b} \ . \tag{14}$$

The formula (14) says that the robot and the obstacle are safely apart. In this case, the invariant coincides with condition (11), which describes the stopping distance.

We prove that the property (13) holds for all executions of Model 1 under the assumption that we start in a state satisfying the following conditions:[8]

$$\phi_{\text{ss}} \;\equiv\; v_r = 0 \land r_c \neq 0 \land \|d_r\| = 1 \;\; . \tag{15}$$

The first condition of the conjunction formalizes that the robot is stopped initially. Note, that any other formula ϕ_{ss} implying the invariant φ_{ss} is a valid starting condition as well. The second conjunct states that the robot is not spinning initially. The last conjunct $\|d_r\| = 1$ says that the direction d_r is a unit vector.

Theorem 1 (Static safety). If the robot starts in a state where ϕ_{ss} (15) holds, then the control model dw_{ss} (Model 1) always guarantees the static safety condition ψ_{ss} (13), as expressed by the provable dL formula

$$\phi_{\text{ss}} \rightarrow [dw_{\text{ss}}]\psi_{\text{ss}} \;\; .$$

We proved Theorem 1 for circular trajectories using KeYmaera [34], a theorem prover for hybrid systems. The proof uses the invariant φ_{ss} (14) for handling the loop. It uses *differential invariants* (16)–(20)—an induction principle for differential equations [31]—to prove properties about *dyn* without relying on symbolic solutions.

$$\begin{aligned}
\eta_{\text{ss}} \equiv\; & t \geq 0 && (16)\\
\land\; & \|d_r\| = 1 && (17)\\
\land\; & v_r = old(v_r) + a_r t && (18)\\
\land\; & -t\left(v_r - \frac{a_r}{2}t\right) \leq p_r^x - old(p_r^x) \leq t\left(v_r - \frac{a_r}{2}t\right) && (19)\\
\land\; & -t\left(v_r - \frac{a_r}{2}t\right) \leq p_r^y - old(p_r^y) \leq t\left(v_r - \frac{a_r}{2}t\right) && (20)
\end{aligned}$$

The differential invariants capture that time progresses (16), that the orientation stays a unit vector (17), that the new speed v_r is determined by the previous speed $old(v_r)$[9] and the acceleration a_r (18), and that the robot does not leave the bounding square of half side length $t(v_r - \frac{a_r}{2}t)$ around its previous position $old(p_r)$ (19)–(20).

6.2 Passive Safety with Maximum Acceleration

Passive safety considers the robot safe if it is able to come to a full stop before making contact with an obstacle. When every moving robot and obstacle follows passive safety then there will be no collisions. Otherwise, i. e., if careless or malicious obstacles are moving in the environment, passive safety ensures that at least our own robot is stopped so that collision impact is kept small. In this section, we will develop a robot controller that provably ensures passive safety. In this section, we remove the restriction that obstacles cannot move, but keep the remaining simplifying assumptions from the beginning ofSection 6:

- in its decisions, the robot will assume it uses maximum braking or maximum acceleration,
- the robot will not be able to reverse its direction, but only drive smooth curves in forward direction,

[8] The formal proof uses the parameter constraints stated earlier, $A \geq 0$, $V \geq 0$, $\Omega \geq 0$, $b > 0$, and $\varepsilon > 0$, which we leave out for simplicity.

[9] The function $old(\cdot)$ is shorthand notation for an auxiliary variable that is initialized to the value of $\cdot$ before the ODE.

Model 2 Dynamic window with passive safety

$$dw_{\text{ps}} \equiv (ctrl_o; ctrl_r; dyn)^* \tag{21}$$
$$ctrl_o \equiv v_o := (*,*);\ ?\|v_o\| \leq V \tag{22}$$
$$ctrl_r \equiv \quad (a_r := -b) \tag{23}$$
$$\cup\ (?v_r = 0;\ a_r := 0;\ \omega_r := 0) \tag{24}$$
$$\cup\ (a_r := A;\ \omega_r := *;\ ?{-\Omega} \leq \omega_r \leq \Omega;\ r_c := *;\ p_o := (*,*);\ ?curve \wedge safe) \tag{25}$$
$$curve \equiv r_c \neq 0 \wedge r_c\omega_r = v_r \tag{26}$$
$$safe \equiv \|p_r - p_o\|_\infty > \frac{v_r^2}{2b} + V\frac{v_r}{b} + \left(\frac{A}{b}+1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon(v_r+V)\right) \tag{27}$$
$$dyn \equiv t := 0;\ \{t' = 1,\ p_o' = v_o,,\ p_r' = v_r d_r,\ v_r' = a_r,\ d_r' = \omega_r d_r^\perp,\ \omega_r' = \frac{a_r}{r_c}\ \&\ v_r \geq 0 \wedge t \leq \varepsilon\} \tag{28}$$

- the robot will not keep track of the center of the circle around which its current arc is taking it, but choose steering through picking a curve radius, and
- obstacles can move, but the robot and the obstacle will decide on their next maneuver at the same time.

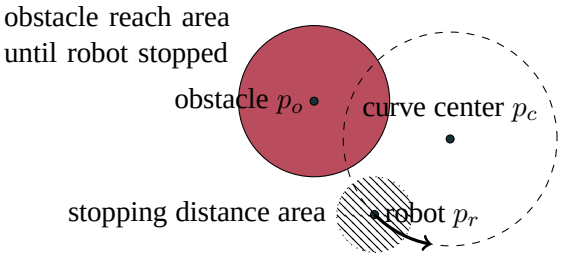

Figure 4: Illustration of passive safety: the area reachable by the robot until it can stop must not overlap with the area reachable by the obstacle during that time.

Modeling We develop a model of the principles in the dynamic window algorithm as a hybrid program, and express passive safety as a safety property in dL. As stated above, the dynamic window algorithm uses the distance to the nearest obstacle for every possible curve to determine admissible velocities. In the presence of moving obstacles, however, all obstacles must be considered and tested for safety (e. g., in a loop). Again, our model exploits the power of nondeterminism to model this concisely by picking *any* obstacle $p_o := (*,*)$ and testing its safety. In each controller run of the robot, the position p_o is updated nondeterministically (again to consider any obstacle including the ones that are now closest because the robot as well as the obstacles moved). In this process, the robot may or may not find another safe trajectory. If it does, the robot can follow that new safe trajectory w.r.t. any nondeterministically chosen obstacle (again, V ensures that all other obstacles will stay more distant than the worst case of the nearest). If not, the robot can still brake on the previous trajectory, which was previously shown to be safe.

Model 2 follows a setup similar to Model 1. The continuous dynamics of the robot and the obstacle as presented in Section 5 above are defined in (28) of Model 2. Model 2 still focuses on circular arcs of non-zero radius. We have made proofs for corresponding dynamics for spinning behavior ($r_c = 0$, $\omega_r \neq 0$) and Archimedean spiral ($\omega_r' = 0$, $a_r \neq 0$) available with KeYmaera as well, but do not discuss them here.

The control of the robot is executed after the control of the obstacle, cf. (21). Note, that both robot and obstacle only write to variables that are read in the plant, but not in the controller of the respective other agent. Therefore, we could swap the controllers to $\mathit{ctrl}_r; \mathit{ctrl}_o$, or even more liberal use a nondeterministic choice of $(\mathit{ctrl}_o; \mathit{ctrl}_r) \cup (\mathit{ctrl}_r; \mathit{ctrl}_o)$, which resembles parallel execution of controllers (see [21] for more details on this technique). Deciding on one specific ordering (here: $\mathit{ctrl}_o; \mathit{ctrl}_r$) reduces proof effort.

The obstacle itself may choose any velocity in any direction up to the maximum velocity V assumed about obstacles ($\|v_o\| \leq V$), cf. (22). This uses the modeling pattern from Section 3: we assign an arbitrary (two-dimensional) value to the obstacle's velocity ($v_o := (*, *)$), which is then restricted to any value up to the maximum velocity using a subsequent test ($?\|v_o\| \leq V$). Overall, (22) allows obstacles to choose an arbitrary velocity in any direction, but at most as fast as V.

The robot follows the same three control choices as in Model 1, see (23)–(25). The main difference to Model 1 is the condition (27), which now has to account for the fact that obstacles may move while the robot tries to avoid collision.

Identification of Safe Controls Based on this shape of the model, again the most critical element is the formula *safe* that we identify as the condition that control choices need to satisfy in order to always keep the robot safe. Here, we extend the intuitive explanation from 6.1 on page 15 to account for the additional obstacle terms in (27).

A circular trajectory of radius r_c ensures passive safety if it allows the robot to stop before it collides with the nearest obstacle. We again consider the extreme case where the radius $r_c = \infty$ is infinitely large and the robot, thus, travels on a straight line. In addition to the stopping distance $\frac{v_r^2}{2B}$ (11) and the distance for compensating acceleration $\left(\frac{A}{b} + 1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon v_r\right)$ (12), the distance between the robot's current position p_r and the nearest obstacle p_o must account for the motion of the obstacle. During the stopping time ($\varepsilon + \frac{v_r + A\varepsilon}{b}$) entailed by (11) and (12), the obstacle might approach the robot on a straight line with maximum velocity V to the point of collision:

$$V\left(\varepsilon + \frac{v_r + A\varepsilon}{b}\right) = V\left(\frac{v_r}{b} + \left(\frac{A}{b} + 1\right)\varepsilon\right) \quad . \tag{29}$$

The safety distance chosen for *safe* in (27) of Model 2 is the sum of the distances (11), (12), and (29). The safety proof will have to show that this construction was safe and that it is also safe for all other curved trajectories that the obstacle and robot could be taking instead.

Verification We verify the safety of the control algorithm modeled as a hybrid program in Model 2, using a formal proof calculus for dL [23, 25, 30]. The robot is safe, if it maintains positive distance $\|p_r - p_o\| > 0$ to the obstacle while the robot is driving (see Table 4):

$$\psi_{\text{ps}} \;\equiv\; v_r \neq 0 \rightarrow (\|p_r - p_o\| > 0) \quad . \tag{30}$$

In order to guarantee ψ_{ps}, the robot must stay at a safe distance, which still allows the robot to brake to a complete stop before the approaching obstacle reaches the robot. The following condition captures this

requirement as an invariant φ_{ps} that we prove to hold for all loop executions:

$$\varphi_{\text{ps}} \;\equiv\; v_r \neq 0 \rightarrow \left(\|p_r - p_o\| > \frac{v_r^2}{2b} + V\frac{v_r}{b} \right) \quad . \tag{31}$$

Formula (31) says that, while the robot is driving, the positions of the robot and the obstacle are safely apart. This accounts for the robot's braking distance $\frac{v_r^2}{2b}$ while the obstacle is allowed to approach the robot with its fastest travel distance $V\frac{v_r}{b}$. We prove that the property (30) holds for all executions of Model 2 under the assumption that we start in a state satisfying the following conditions:[10]

$$\phi_{\text{ps}} \;\equiv\; v_r = 0 \land r_c \neq 0 \land \|d_r\| = 1 \quad . \tag{32}$$

The first condition of the conjunction formalizes that the robot is stopped initially, whereas the second conjunct states that the robot is not spinning. The last conjunct $\|d_r\| = 1$ says that the direction d_r is a unit vector. Again, any other condition ϕ_{ps} such that $\phi_{\text{ps}} \rightarrow \varphi_{\text{ps}}$ holds would work as well.

Theorem 2 (Passive safety). If the robot starts in a state where ϕ_{ps} (32) holds, then the control model dw_{ps} (Model 2) always guarantees the passive safety condition ψ_{ps} (30), as expressed by the provable $\mathsf{d}\mathcal{L}$ formula

$$\phi_{\text{ps}} \rightarrow [dw_{\text{ps}}]\psi_{\text{ps}} \quad .$$

We proved Theorem 2 for circular trajectories, spinning, and spiral trajectories using KeYmaera [34], a theorem prover for hybrid systems. The proof uses the invariant φ_{ps} (31). It extends the differential invariants (16)–(20) for static safety with invariants about obstacle motion: similar to the robot, the obstacle does not leave its bounding square of half side length tV around its previous position $old(p_o)$ (33)–(34)

$$\eta_{\text{ps}} \equiv \eta_{\text{ss}} \land -tV \leq p_o^x - old(p_o^x) \leq tV \tag{33}$$

$$\land -tV \leq p_o^y - old(p_o^y) \leq tV \tag{34}$$

In the next section, we explore the stronger requirements of passive friendly safety, where the robot not only stops safely itself, but also allows for the obstacle to stop before a collision occurs.

6.3 Passive Friendly Safety of Obstacle Avoidance

Passive friendly safety, as introduced above, requires the robot to take careful decisions that respect the features of moving obstacles. The intuition behind passive friendly safety is that our own robot should retain space for careless obstacles. Passive friendly safety ensures that there will not be collisions, as long as every obstacle makes an effort of stopping to avoid collision when it sees the robot, even when some obstacles approach intersections carelessly and turn around corners without looking first. The definition of Maček et al. [17] requires that the robot respects the worst-case braking time of the obstacle. In our model, the worst-case braking time is represented as follows. We assume an upper bound τ on the reaction time of the obstacle and a lower bound b_o on its braking capabilities. Then, τV is the maximal distance that the obstacle can travel before actually reacting and $\frac{V^2}{2b_o}$ is the distance for the obstacle to stop from the maximal velocity V with an assumed minimum braking capability b_o.

[10] The formal proof uses the parameter constraints stated earlier, $v_r \geq 0$, $A \geq 0$, $V \geq 0$, $\Omega \geq 0$, $b > 0$, and $\varepsilon > 0$, which we leave out for simplicity.

Modeling Model 3 uses the same obstacle avoidance algorithm as Model 2. The essential difference reflects what the robot considers to be a safe distance to an obstacle. As shown in (36) the distance not only accounts for the robot's own braking distance, but also for the braking distance $\frac{V^2}{2b_o}$ and reaction time τ of the obstacle. The verification of passive friendly safety is more complicated than passive safety as it accounts for the behavior of the obstacle discussed below.

Model 3 Dynamic window with passive friendly safety

$$dw_{\text{pfs}} \equiv (\mathit{ctrl}; \mathit{dyn})^* \text{ (see Model 2 for details)} \tag{35}$$

$$\mathit{safe} \equiv \|p_r - p_o\|_\infty > \frac{v_r^2}{2b} + V\frac{v_r}{b} + \frac{V^2}{2b_o} + \tau V + \left(\frac{A}{b} + 1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon(v_r + V)\right) \tag{36}$$

Verification We verify the safety of the robot's control choices as modeled in Model 3. Unlike the passive safety case, the passive friendly safety property ϕ_{pfs} should guarantee that if the robot stops, moving obstacles (cf. Model 4) still have enough time and space to avoid a collision. This requirement can be captured by the following dL formula:

$$\eta_{\text{pfs}} \;\equiv\; \big((v_r = 0) \wedge \eta_{\text{obs}} \wedge (0 \le v_o \le V)\big) \to \langle \mathit{obstacle} \rangle \, (\|p_r - p_o\| > 0 \wedge v_o = 0) \tag{37}$$

where the property η_{obs} accounts for the stopping distance of the obstacle:

$$\eta_{\text{obs}} \;\equiv\; \|p_r - p_o\| > \frac{V^2}{2b_o} + \tau V \;\; .$$

Formula (37) says that, once the robot stops ($v_r = 0$), there exists an execution of the hybrid program *obstacle*, (existence of a run is formalized by the diamond operator $\langle \mathit{obstacle} \rangle$), that allows the obstacle to stop ($v_o = 0$) without colliding ($\|p_r - p_o\| > 0$). Passive friendly safety ψ_{pfs} is now stated as

$$\psi_{\text{pfs}} \;\equiv\; \eta_{\text{obs}} \wedge \eta_{\text{pfs}} \;\; .$$

We study passive friendly safety with respect to the initial feasible states satisfying the following property:

$$\phi_{\text{pfs}} \;\equiv\; \eta_{obs} \wedge r_c \neq 0 \wedge \|d_r\| = 1 \;\; . \tag{38}$$

Observe that, in addition to the condition η_{pfs}, the difference with the passive safety condition is reflected in the special treatment of the case $v_r = 0$. Indeed, if the robot would start with translational velocity $v_r = 0$ (which is passive safety) while not satisfying η_{obs}, then we cannot prove passive friendly safety as the obstacle may be unsafely close already initially. Besides, we can see in ψ_{pfs} that we are required to prove η_{pfs} even when the robot comes to a full stop.

In Model 3 the hybrid program $ctrl_o$ is a coarse model given by equation (22), which only constrains its non-negative velocity to be less than or equal to V. Such an obstacle could trivially prevent a collision by stopping instantaneously, and is, thus, a trivial obstacle model to consider for passive friendly safety. In η_{pfs}, we, instead, consider a more interesting refined obstacle behavior modeled by the hybrid program given in Model 4.

The refined obstacle may choose any direction as described by the unit vector d_o in (40) and any acceleration a_o, as long as it does not exceed the velocity bound V. The dynamics of the obstacle are straight ideal-world translational motion in the two-dimensional plane, see (41).

Model 4 Refined obstacle with acceleration control

$$\mathit{obstacle} \equiv (\mathit{ctrl}_{\tilde{o}};\ \mathit{dyn}_{\tilde{o}})^* \tag{39}$$

$$\mathit{ctrl}_{\tilde{o}} \equiv d_o := (*,*);\ ?\|d_o\| = 1;\ a_o := *;\ ?v_o + a_o\varepsilon_o \leq V \tag{40}$$

$$\mathit{dyn}_{\tilde{o}} \equiv t := 0;\ \{t' = 1,\ p_o' = v_o d_o, v_o' = a_o \ \&\ t \leq \varepsilon_o \wedge v_o \geq 0\} \tag{41}$$

Theorem 3 (Passive friendly safety). If the property ϕ_{pfs} (38) holds initially, then the control model dw_{pfs} (Model 3), guarantees the passive friendly safety condition ψ_{pfs} in presence of obstacles per Model 4, as expressed by the provable $\mathsf{d}\mathcal{L}$ formula

$$\phi_{\text{pfs}} \rightarrow [dw_{\text{pfs}}]\psi_{\text{pfs}} \ .$$

We verified Theorem 3 in KeYmaera. The symbolic bounds on velocity, acceleration, braking, and time in the above models represent uncertainty implicitly (e. g., the braking power b can be instantiated with the minimum specification of the robot's brakes, or with the actual braking power achievable w.r.t. the current terrain). Dually, whenever knowledge about the current state is available, the bounds can be instantiated more aggressively to allow efficient robot behavior. For example, in a rare worst case we may face a particularly fast obstacle, but right now there are only slow-moving obstacles around. Theorems 2–3 are verified for all those values. Section 7.5 illustrates how to explicitly model different kinds of obstacles in a single model. Other aspects of uncertainty need explicit changes in the models and proofs, as discussed in the next section.

6.4 Passive Orientation Safety

So far, we did not consider orientation as part of the safety specification. The notion of passive safety requires the robot to stop to avoid imminent collision, which can be inefficient or even impossible when sensor coverage is not exhaustive. For example, if an obstacle is close behind the robot (cf. Fig. 5a), the robot would have to stop to guarantee passive safety, while it could choose a new curve that leads away from the obstacle with a more liberal safety notion.

We introduce *passive orientation safety* that only requires the robot to remain safe with respect to the obstacles in its orientation of responsibility. Overall system safety depends on the sensor coverage of the robot and the obstacles. For example, if two robots drive side-by-side with only very narrow sensor coverage to the front, they still might collide when their paths cross. Even with limited sensor coverage, if both robots can observe some separation markers in space (e. g., lane markers) that keeps their paths separated, then passive orientation safety ensures that there will not be collisions. Likewise, passive orientation safety ensures that there will not be collisions when every robot and obstacle covers $180°$ in its orientation of responsibility, i. e., it can observe its vicinity to the front, but not to the rear.

This notion of safety is suitable for structured spaces where obstacles can easily determine the trajectory and observable region of the robot (e. g., lanes on streets). The robot is responsible for collisions inside its observable area ("field of vision", cf. Fig. 5b) and has to ensure that it can stop if needed before leaving the observable region (cf. Fig. 5c), because it could otherwise cause collisions when moving into the blind spot just outside its observable area. However, the robot does not make guarantees for obstacles that it cannot see. If an obstacle starts outside the observable region and subsequently hits the robot, then it is considered the fault of the obstacle. If the robot guarantees passive orientation safety and every obstacle outside the observable region guarantees that it will not interfere with the robot, a collision between the robot and an

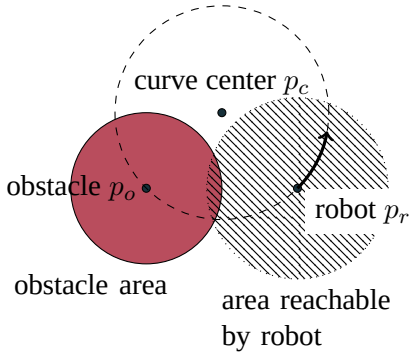

(a) Passive safety without orientation: the robot cannot choose a new curve if both the robot area and the trajectory overlap with the obstacle area, even if that curve leads away from the obstacle behind.

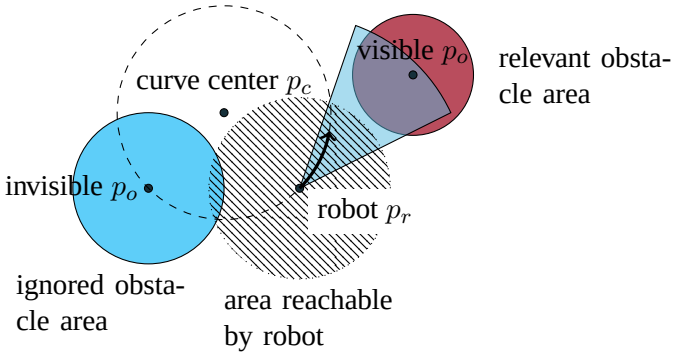

(b) The area observable by the robot (circular sector centered at robot): the distance to all visible obstacles must be safe. Obstacles outside the visible area are responsible for avoiding collision with the robot until they become visible, i. e., obstacles are assumed to not blindside the robot.

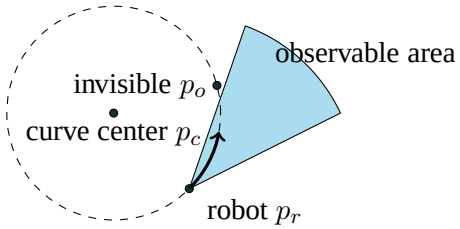

(c) An obstacle might sit on the robot's curve just outside the observable area. Thus, the robot must ensure that it can stop inside its current observable area, since starting to brake when the obstacle becomes visible might be too late.

Figure 5: Passive orientation safety compared to passive safety without orientation

obstacle never happens while the robot is moving. In fact, collisions can be avoided when obstacles do not cross the trajectory of the robot. Any obstacles inside the observable region can drive with passive safety restrictions (i. e., guarantee not to exceed a maximum velocity) because the robot will brake or choose a new curve to avoid collisions. Obstacles that start outside the observable region can rely on the robot to stay inside its observable region, and that it guarantees collision avoidance for the obstacles it can see. Also, it guarantees to only enter places it can see (i. e. the robot will stop before it drives to places that it did not see when evaluating the safety of a curve).

Modeling To express that an obstacle was invisible to the robot when it chose a new curve we introduce the non-rigid function *Visible*, with $\mathit{Visible} > 0$ indicating that the obstacle was visible to the robot when it chose a new curve. The observable region is aligned with the orientation of the robot and extends symmetrically by a constant design parameter γ that captures the angular width of the field of vision: The robot can see everything within angle $\frac{\gamma}{2}$ to its left or right. With the newly introduced notions, the safety condition to guarantee passive orientation safety can be expressed in QdL as follows:

$$\psi_{\text{pos}} \equiv v_r \neq 0 \rightarrow \big(\|p_r - p_o\| > 0 \vee (\mathit{Visible} \leq 0 \wedge |\beta| < \gamma)\big) \tag{42}$$

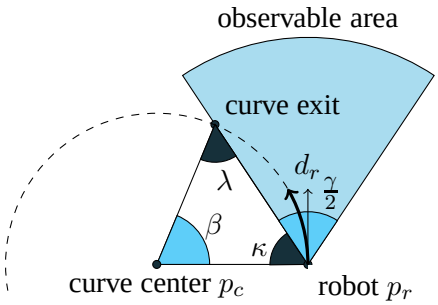

Figure 6: Determining the point where the curve exits the observable region of angular width γ by keeping track of the angular progress β along the curve: $\kappa = 90° - \frac{\gamma}{2}$ because γ extends equally to both sides of the orientation d_r, which is perpendicular to the line from the robot to p_c (because d_r is tangential to the curve). $\lambda = \kappa$ because the triangle is isosceles. Thus, $\beta = 180° - \kappa - \lambda = \gamma$ at exactly the moment when the robot would leave the observable region.

This means that, when the robot is driving, every obstacle is either sufficiently far away or it came from outside the observable region while the robot stayed inside. For determining whether or not the robot stayed inside the observable region, we compare the angular progress β along the curve with the angular width γ of the observable region, see Fig. 6 for details. The angular progress β is reset to zero each time the robot chooses a new curve in (44) and evolves according to $\beta' = \omega_r$ when the robot moves (50). Thus, β always holds the value of the angle on the current curve between the current position of the robot and its position when it chose the curve. Note that passive safety is a special case of passive orientation safety for $\gamma = \infty$.[11]

The new robot controller now only takes obstacles in its observable region into account when computing the safety of a new curve in *safe* (48). It chooses curves such that it can stop before leaving the observable region, i. e. it ensures a clear distance ahead (*cda*): such a curve is characterized by the braking distance of the robot being less than $\gamma \cdot |r|$, which is the length of the arc between the starting position when choosing the curve and the position where the robot would leave the observable region, cf. Fig. 6. When the robot successfully chooses a new curve (44), the status of the obstacle (i. e. whether or not it is visible) is stored in *Visible* so that the visibility state is available when checking the safety property.

Verification We formalize passive orientation safety in Theorem 4.

Theorem 4 (Passive orientation safety). When started in an initially passive-orientation-safe state, the model dw_{pos} (Model 5) ensures passive orientation safety:

$$\phi_{\text{pos}} \rightarrow [dw_{\text{pos}}]\psi_{\text{pos}} \ .$$

We proved Theorem 4, i. e., that Model 5 is safe with respect to passive orientation safety.

7 Refined Models for Safety Verification

So far, the models used for safety verification made several simplifying assumptions to focus on the basics of different safety notions. In this section, we discuss how to create more realistic models with actual

[11] Passive orientation safety restricts admissible curves to those where the robot can stop before $|\beta| > \gamma$ and the model does not take advantage of the fact that $360°$ subsumes unrestricted visibility.

Model 5 Passive orientation safety

$$dw_{\text{pos}} \equiv (ctrl_o; ctrl_r; dyn)^* \tag{43}$$

$$\begin{aligned} ctrl_r \equiv\ & (a_r := -b) \\ & \cup (?v_r = 0;\ a_r := 0;\ \omega_r := 0) \\ & \cup (a_r := A;\ \omega_r := *;\ ? -\Omega \leq \omega_r \leq \Omega; \end{aligned} \tag{44}$$

$$r_c := *;\ p_o := (*, *);\ Visible := *; \tag{45}$$

$$?curve \wedge safe \wedge cda); \tag{46}$$

$$curve \equiv r_c \neq 0 \wedge r_c\omega_r = v_r \tag{47}$$

$$safe \equiv Visible > 0 \rightarrow \|p_r - p_o\|_\infty > \frac{v_r^2}{2b} + V\frac{v_r}{b} + \left(\frac{A}{b} + 1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon(v_r + V)\right) \tag{48}$$

$$cda \equiv \gamma|r| > \frac{v_r^2}{2b} + \left(\frac{A}{b} + 1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon v_r\right) \tag{49}$$

$$dyn \equiv t := 0;\ \beta := 0; \tag{50}$$

$$\{p_r' = v_r d_r,\ d_r' = \omega_r d_r^\perp,\ v_r' = a_r,\ \beta' = \omega_r,\ \omega_r' = \frac{a_r}{r},\ p_o' = v_o,\ t' = 1\ \&\ v_r \geq 0 \wedge t \leq \varepsilon\} \tag{51}$$

acceleration, distance measurement along the trajectory, measurement uncertainty and actuator disturbance, non-synchronized control of obstacle and robot, and explicit representation of arbitrary many obstacles. We introduce the model extensions for passive safety as an example. The extensions apply to the other safety notions in a similar fashion.

7.1 Passive Safety with Actual Acceleration

Model 2 uses the robot's maximum acceleration in its safety requirement (27) when it determines whether or not a new curve will be safe. This condition is conservative, since the robot can only decide between maximum acceleration or maximum braking. If (27) does not hold (which is independent from the chosen curve, i. e. the radius r_c), then Model 2 forces a driving robot to brake with maximum deceleration $-b$, even if it might be safe to just slightly brake or just not accelerate in full. As a result, Model 2 lacks efficiency in that it may take the robot longer to reach a goal because it has to make extreme choices. Besides efficiency concerns, extreme choices can be undesirable for comfort reasons (e. g., decelerating a car braking with full power should be reserved for emergency cases).

Fig. 7 illustrates how safety constraint (27) represents the maximally conservative choice: it forces the robot to brake (the outermost circle around the robot p_r intersects with the obstacle), even though many points reachable with $-b \leq a_r < A$ are still safe (solid blue area does not intersect with the obstacle).

Modeling Model 6 adapts Model 2 to actual acceleration, i. e., in (56) the robot chooses any acceleration in the physical limits $-b \leq a_r \leq A$ instead of just maximum acceleration.

This change requires us to adapt the control condition (59) that keeps the robot safe. The next paragraph first gives the intuition behind condition (59), before in 7.1 we ultimately justify its correctness with a safety proof.

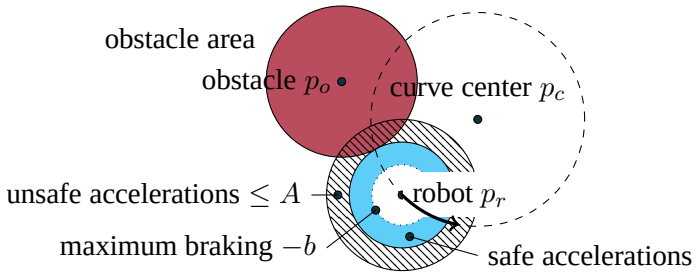

Figure 7: Passive safety with actual acceleration: the actual acceleration choice $-b \leq a_r \leq A$ must not take the robot into the area reachable by the obstacle. Dotted circle around robot position p_r: earliest possible stop with maximum braking $-b$; solid blue area between dotted circle and dashed area: safe a_r; dashed area: reachable with unsafe accelerations.

Model 6 Passive safety with actual acceleration

$$dw_{\text{psa}} \equiv (ctrl_o; ctrl_r; dyn)^* \tag{52}$$
$$ctrl_o \equiv \text{ see Model 2} \tag{53}$$
$$ctrl_r \equiv \quad (a_r := -b) \tag{54}$$
$$\cup\ (?v_r = 0;\ a_r := 0;\ \omega_r := 0) \tag{55}$$
$$\cup\ (a_r := *;\ ? -b \leq a_r \leq A;\ \omega_r := *;\ ?{-}\Omega \leq \omega_r \leq \Omega; \tag{56}$$
$$r_c := *;\ p_o := (*,*);\ ?curve \wedge safe) \tag{57}$$
$$curve \equiv \text{ see Model 2} \tag{58}$$
$$safe \equiv \|p_r - p_o\|_\infty > \begin{cases} dist_{\geq} & \text{if } v_r + a_r\varepsilon \geq 0 \\ dist_{<} & \text{otherwise} \end{cases} \tag{59}$$
$$dyn \equiv \text{ see Model 2} \tag{60}$$

Identification of Safe Constraints Here, we follow [16] to relax constraint (27) so that the robot can choose any acceleration $-b \leq a_r \leq A$ and checks this actual acceleration for safety. That way, it only has to fall back to the emergency braking branch $a_r := -b$ if there is no other safe choice available. We distinguish two cases:

- $v_r + a_r\varepsilon \geq 0$: All acceleration choices $-b \leq a_r \leq A$ keep a nonnegative velocity if followed for the full cycle duration ε.
- $v_r + a_r\varepsilon < 0$: Some acceleration choice $-b \leq a_r < 0$ cannot be followed for the full duration ε without stopping the evolution to prevent a negative velocity.

In the first case we continue to use formula (27) with actual a_r substituted for A to compute the safety distance:

$$dist_{\geq} = \frac{v_r^2}{2b} + V\frac{v_r}{b} + \left(\frac{a_r}{b} + 1\right)\left(\frac{a_r}{2}\varepsilon^2 + \varepsilon(v_r + V)\right) \tag{61}$$

In the second case, distance (61) is unsafe, because the terminal velocity when following a_r for ε time is negative (unlike case 1). Thus, the robot may have collided at a time before ε, while the term in (61) only indicates that it will no longer be in a collision state at time ε. Consider the time t_b when the robot's velocity becomes zero ($v_r + a_r t_b = 0$) so that its physics stops (recall that braking does not make the robot move backwards but merely stop). Hence, $t_b = -\frac{v_r}{a_r}$ since the first case applies for $a_r = 0$. In duration t_b the robot will drive a total distance of $dist_r = -\frac{v_r^2}{2a_r} = \int_0^{t_b}(v_r + a_r t)dt$. The obstacle may drive up to $dist_o = Vt_b$ until the robot is stopped. Thus, we compute distance using (62) to account for the worst case that both robot and obstacle drive directly towards each other (note that $-b \leq a_r < 0$).

$$\mathit{dist}_< = -\frac{v_r^2}{2a_r} - V\frac{v_r}{a_r} \tag{62}$$

Verification We verify the safety of the control algorithm as modeled in Model 6 in the dL proof calculus.

Theorem 5 (Passive safety with actual acceleration). If the robot starts in a state where ϕ_{ps} (32) holds, then the control model dw_{psa} (Model 6) guarantees the passive safety condition ψ_{ps} (30), as expressed by the provable dL formula

$$\phi_{\text{ps}} \rightarrow [dw_{\text{psa}}]\psi_{\text{ps}} \ .$$

We proved Theorem 5 in KeYmaera.

7.2 Passive Safety with Trajectory Distance Measurement

Models 2 and 6 used a safety distance in infinity norm for the safety constraints, which grossly overapproximates the actual trajectory of the robot by a box around the robot. For example, recall the safety distance (27) of Model 2

$$\|p_r - p_o\|_\infty > \frac{v_r^2}{2b} + V\frac{v_r}{b} + \left(\frac{A}{b} + 1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon(v_r + V)\right) \ ,$$

which needs to be large enough in either one axis, irrespective of the actual trajectory that the robot will be taking. This constraint can be inefficient when the robot chooses a trajectory that will keep it close to its current position (e. g., when driving along a small circle). For example, a robot with constant velocity $v_r = 4$ and reaction time $\varepsilon = 1$ will traverse a small circle with radius $r = \frac{1}{\pi}$ and corresponding circumference $2\pi r = 2$ twice within time ε. So far, the safety constraint required the total distance of 4 as a safety distance between the robot and the obstacle, because it overapproximated its actual trajectory by a box. However, the robot actually never moves away more than $\frac{2}{\pi}$ from its original position because it moves on a circle (cf. Fig. 8a).

Modeling We change the robot controller to improve its efficiency. One choice would be to explicitly express circular motion in terms of sine and cosine and then compute all possible positions of the robot explicitly. However, besides being vastly inefficient in a real controller, this introduces transcendental functions and would leave the decidable part of the arithmetic of real-closed fields [25]. Hence, we will use the distance of the obstacle to the trajectory itself in the control conditions. Such a distance computation requires that we adapt the constraint *curve* (70) to express the curve center explicitly. So far, the curve was uniquely determined by the radius r_c and the orientation d_r of the robot. Now that we need the curve center explicitly for distance calculation to the obstacle, we have to choose the curve center p_c such that

- $(p_r - p_c)$ is perpendicular to the robot orientation d_r, i. e., d_r is tangential to the curve, and

- located correctly to the left or right of the robot, so that it fits to the clockwise or counter-clockwise motion as signified by r_c.

Thus, the condition *curve* (70) now checks if the chosen curve and the direction of the robot are consistent, i. e., $|r_c| = \|p_r - p_c\|$ and $d_r = \frac{(p_r - p_c)^\perp}{r_c}$ must hold. Additionally, we augment the robot with a capability to turn around in place when stopped ($v_r = 0$). For this, (67) is extended with a choice of either turning around ($d_r := -d_r$) or remaining oriented as is ($d_r := d_r$), and the corresponding choice of a curve center p_c such that the curve variables remain consistent according to the subsequent test $?\mathit{curve}$.

Identification of Safe Controls With the changes in distance measurement introduced above, we now relax the control conditions that keep the robot safe. The distance of the obstacle to the trajectory can be described with the following two components:

1. Calculate the distance of the obstacle to the circle: $||r_c| - \|p_o - p_c\||$, which is the distance between the obstacle and the circle center minus the radius.

2. Calculate the maximum distance that the obstacle can drive until the robot comes to a stop. This distance is equal to the distances calculated in the previous models, i. e. in the case $v_r + a_r\varepsilon \leq 0$ it is $-V\frac{v_r}{a_r}$ and in the case $v_r + a_r\varepsilon \geq 0$ it is $V\left(\varepsilon + \frac{v_r + a_r\varepsilon}{b}\right)$.

If the distance between the obstacle and the circle describing the robot's trajectory is greater than the sum of those distances then the robot can stop before hitting the obstacle and, thus, choosing the new curve is safe, which leads us to choose the following safety condition:

$$||r_c| - \|p_o - p_c\|| > \begin{cases} V\left(\varepsilon + \frac{v_r + a_r\varepsilon}{b}\right) & \text{if } v_r + a_r\varepsilon \geq 0 \\ -V\frac{v_r}{a_r} & \text{otherwise} \end{cases} \tag{63}$$

We use condition (63), which now uses the Euclidean norm, for choosing a new curve in Model 7. With this new constraint, the robot is allowed to choose the curve in Fig. 8a. However, constraint (63) has drawbacks when the trajectory of the robot is slow along a large circle and the obstacle is close to the circle, as illustrated in Fig. 8b. In this case the robot is only allowed to choose very small accelerations because the obstacle is very close to the circle. A disjunction of constraints (59) and (63) provides the best of both worlds, see (71).

Verification We verify the safety of the robot's control choices as modeled in Model 7 in the $\mathsf{d}\mathcal{L}$ proof calculus.

Theorem 6 (Passive safety with trajectory distance measurement). If the property ϕ_{ps} holds initially, then the model $\mathit{dw}_{\text{psdm}}$ (Model 7) guarantees passive safety, as expressed by the provable $\mathsf{d}\mathcal{L}$ formula

$$\phi_{\text{ps}} \rightarrow [\mathit{dw}_{\text{psdm}}]\psi_{\text{ps}} \ .$$

We proved Theorem 6 with KeYmaera.

7.3 Passive Safety Despite Uncertainty

Robots have to deal with uncertainty in almost every aspect of their interaction with the environment, ranging from sensor inputs (e. g., inaccurate localization, distance measurement) to actuator effects (e. g., wheel slip

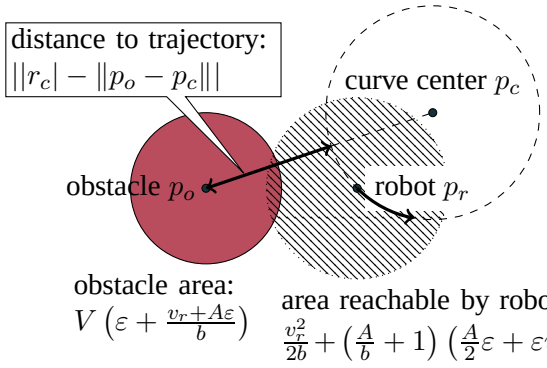

(a) Safe since the obstacle area does not overlap the dashed trajectory.

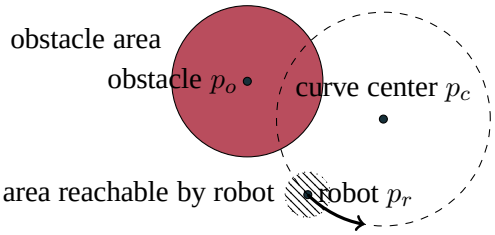

(b) Safe since obstacle area and dotted robot area do not overlap.

Figure 8: Safe robot trajectories

depending on the terrain). In this section, we show how the three most important classes of uncertainty can be handled explicitly in the models. First, we allow localization uncertainty, so that the robot knows its position only approximately, which has a considerable impact on uncertainty over time. We then consider imperfect actuator commands, which means that the effective physical braking and acceleration will differ from the controller's desired output. Finally, we allow velocity uncertainty, so the robot knows its velocity only approximately. We use nondeterministic models of uncertainty as intervals around the real position, acceleration, and velocity, without any probabilistic assumptions about their distribution.[12] Such intervals are instantiated, e. g., according to sensor or actuator specification (e. g., GPS error), or w.r.t. experimental measurements.[13]

7.3.1 Location Uncertainty

Model 8 introduces location uncertainty. It adds a location measurement $\hat{p}_r$ before the control decisions are made such that the controller only bases its decisions on the most recent location measurement $\hat{p}_r$, which can deviate from the true location p_r. This location measurement may deviate from the real position p_r by no more than the symbolic parameter Δ_p, cf. (74). The measured location $\hat{p}_r$ is used in all control decisions of the robot (e. g., in (75) to compute whether or not it is safe to change the curve), while the robot's physical motion is still computed on the symbolic real position p_r.

Theorem 7 (Passive safety despite location uncertainty). If $\phi_{\text{ps}} \wedge \Delta_p \geq 0$ holds initially, then the model dw_{pslu} (Model 8) guarantees passive safety, as expressed by the provable $\mathsf{d}\mathcal{L}$ formula

$$\phi_{\text{ps}} \wedge \Delta_p \geq 0 \rightarrow [dw_{\text{pslu}}]\psi_{\text{ps}} \quad .$$

We proved Theorem 7 in KeYmaera, so that the robot guarantees passive safety if it accounts for the location uncertainty as stated in the safety constraint (75).

[12] Other error models are supported, as long as they are clipped to guaranteed intervals, because in the safety proof we have to analyze all measured values, regardless of their probability. For an advanced analysis technique considering probabilities, see stochastic $\mathsf{d}\mathcal{L}$ [27].

[13] Instantiation with probabilistic bounds means that the symbolically guaranteed safety is traded for a safety probability.

Model 7 Passive safety when considering the trajectory of the robot in distance measurement, extends Model 6

$$\mathit{dw}_{\text{psdm}} \equiv (\mathit{ctrl}_o; \mathit{ctrl}_r; \mathit{dyn})^* \tag{64}$$

$$\mathit{ctrl}_o \equiv \text{see Model 2} \tag{65}$$

$$\mathit{ctrl}_r \equiv \quad (a_r := -b) \tag{66}$$

$$\cup \left(?v_r = 0;\ a_r := 0;\ w_r := 0;\ (d_r := -d_r \cup d_r := d_r);\ r_c := *;\ p_c := (*,*);\ ?\mathit{curve}\right) \tag{67}$$

$$\cup (a_r := *;\ ? -b \leq a_r \leq A;\ \omega_r := *;\ ? -\Omega \leq \omega_r \leq \Omega; \tag{68}$$

$$r_c := *;\ p_c := (*,*);\ p_o := (*,*);\ ?\mathit{curve} \wedge \mathit{safe}) \tag{69}$$

$$\mathit{curve} \equiv r_c \neq 0 \wedge |r_c| = \|p_r - p_c\| \wedge d_r = \frac{(p_r - p_c)^\perp}{r_c} \wedge r_c\omega_r = v_r \tag{70}$$

$$\begin{aligned} \mathit{safe} \equiv \quad & \left(\|p_r - p_o\|_\infty > \begin{cases} \mathit{dist}_{\geq} & \text{if } v_r + a_r\varepsilon \geq 0 \\ \mathit{dist}_{<} & \text{otherwise} \end{cases} \right) \\ \vee & \left(||r_c| - \|p_o - p_c\|| > \begin{cases} V\left(\varepsilon + \frac{v_r + a_r\varepsilon}{b}\right) & \text{if } v_r + a_r\varepsilon \geq 0 \\ -V\frac{v_r}{a_r} & \text{otherwise} \end{cases} \right) \end{aligned} \tag{71}$$

$$\mathit{dyn}_1 \equiv \text{see Model 2} \tag{72}$$

Model 8 Passive safety despite location uncertainty, extends Model 2

$$\mathit{dw}_{\text{pslu}} \equiv (\mathit{locate}; \mathit{ctrl}; \mathit{dyn})^* \text{ (see Model 2 for details)} \tag{73}$$

$$\mathit{locate} \equiv \hat{p}_r := (*,*);\ ?\|\hat{p}_r - p_r\| \leq \Delta_p \tag{74}$$

$$\mathit{safe} \equiv \|\hat{p}_r - p_o\|_\infty > \frac{v_r^2}{2b} + V\frac{v_r}{b} + \Delta_p + \left(\frac{A}{b} + 1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon(v_r + V)\right) \tag{75}$$

7.3.2 Actuator Perturbation

Model 9 introduces actuator perturbation between control and dynamics, cf. (76). Actuator perturbation affects the acceleration by a damping factor δ_a, known to be at most a maximum damping Δ_a, i. e., $\delta_a \in [\Delta_a, 1]$, cf. (77). Note, that this damping factor can change arbitrarily often, but is assumed to be constant during the continuous evolution that takes a maximum of ε time units. The perturbation may cause the robot to have full acceleration ($\delta_a = 1$) but fully reduced braking ($\delta_a = \Delta_a$). For instance, the robot accelerates on perfect terrain, but is unlucky enough to be on slippery terrain again when it brakes. The robot considers this worst-case scenario during control in its safety constraint (78).

Theorem 8 (Passive safety despite actuator perturbation). If $\phi_{\text{ps}} \wedge \Delta_a > 0$ holds initially, then the model $\mathit{dw}_{\text{psap}}$ (Model 9) guarantees passive safety, as expressed by the provable dL formula

$$\phi_{\text{ps}} \wedge \Delta_a > 0 \rightarrow [\mathit{dw}_{\text{psap}}]\psi_{\text{ps}} \ .$$

We proved Theorem 8 in KeYmaera, so that the robot guarantees passive safety if it accounts for the actuator perturbation as stated in the safety constraint (78).

Model 9 Passive safety despite actuator perturbation, extends Model 2

$$\textit{dw}_{\text{psap}} \equiv (\textit{ctrl}; \textit{act}; \textit{dyn})^* \text{ (see Model 2 for details)} \tag{76}$$

$$\textit{act} \equiv \delta_a := *;\ ?(0 < \Delta_a \leq \delta_a \leq 1);\ \tilde{a}_r := \delta_a a_r; \tag{77}$$

$$\textit{safe} \equiv \|p_r - p_o\|_\infty > \frac{v_r^2}{2b\Delta_a} + V\frac{v_r}{b\Delta_a} + \left(\frac{A}{b\Delta_a} + 1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon(v_r + V)\right) \tag{78}$$

dyn of Model 2 with a_r replaced by $\tilde{a}_r$

7.3.3 Velocity Uncertainty

Model 10 Passive safety despite velocity uncertainty, extends Model 2

$$\textit{dw}_{\text{psvu}} \equiv (\textit{ctrl}; \textit{dyn})^* \text{ (see Model 2 for details)} \tag{79}$$

$$\textit{ctrl}_r \equiv \hat{v} := *;\ ?(\hat{v} \geq 0 \wedge v_r - \Delta_v \leq \hat{v} \leq v_r + \Delta_v); \tag{80}$$

$$(a_r := -b)$$

$$\cup\ (?v_r = 0;\ a_r := 0;\ \omega_r := 0) \tag{81}$$

$$\cup\ (a_r := *;\ ? - b \leq a_r \leq A;\ \omega_r := *;\ ?{-}\Omega \leq \omega_r \leq \Omega; \tag{82}$$

$$p_c := (*, *);\ d_r := (*, *);\ p_o := (*, *);\ ?\textit{curve} \wedge \textit{safe})$$

$$\textit{safe} \equiv \|p_r - p_o\|_\infty > \frac{(\hat{v} + \Delta_v)^2}{2b} + V\frac{\hat{v} + \Delta_v}{b} + \left(\frac{A}{b} + 1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon(\hat{v} + \Delta_v + V)\right) \tag{83}$$

Model 10 introduces velocity uncertainty. To account for the uncertainty, the robot reads off a (possibly inexact) measurement $\hat{v}$ of its velocity v_r at the beginning of its control phase (80). It knows that the measured velocity $\hat{v}$ deviates by at most a measurement error Δ_v from the actual velocity v_r. Also, the robot knows that its velocity is non-negative. Thus, we can assume that $\hat{v}$ is always equal to or greater than zero by a corresponding transformation of the measurements. In order to stay safe, the controller has to make sure that the robot stays safe even if its true velocity is maximally larger than the measurement, i. e. $v_r = \hat{v} + \Delta_v$. The idea is now that the controller makes all control choices with respect to the maximal velocity $\hat{v} + \Delta_v$ instead of the actual velocity v_r. The continuous evolution, in contrast, still uses the actual v_r, of course, because the robot's physics will not be confused by a sensor measurement error. Now, since we used the maximal possible velocity when considering the safety of new curves in the controller we can prove that the robot will still be safe. A problem when using $\hat{v}$ instead of v_r occurs in the second branch (80) of $\textit{ctrl}_r$: Because of the velocity uncertainty we no longer know if v_r is zero (i. e. the robot is stopped), which is necessary for changing the direction of the robot. However, one can argue that even with velocity uncertainty the robot should still know when it came to a stop. Thus, we leave the branch unchanged and still use the test $?v_r = 0$ inside it.

Theorem 9 (Passive safety despite velocity uncertainty). If started in a state where ϕ_{ps} holds, then the model $\textit{dw}_{\text{psvu}}$ (Model 10) guarantees passive safety, as expressed by the provable $\mathsf{d}\mathcal{L}$ formula $\phi_{\text{ps}} \wedge \Delta_v \geq 0 \rightarrow [\textit{dw}_{\text{psvu}}]\psi_{\text{ps}}$.

We proved Theorem 9 in KeYmaera.

7.4 Non-Synchronized Control of Obstacle and Robot

In the models so far, the controllers of the robot and the obstacle were executed synchronously, i. e., the robot and the obstacle made their control decisions at the same time. While the obstacle could always choose its previous control choices again if it does not want to act, the previous models only allowed the obstacle to decide when the robot made a control decision, too.[14] This does not reflect reality perfectly, since we want liberal obstacle models without assumptions about when an obstacle makes a control decision. We want to make sure that the robot remains safe regardless of how often and at which times the obstacles change their speed and orientation.

Model 11 Non-Synchronized Control of Obstacle and Robot

$$dw_{\text{psns}} \equiv (ctrl_r;\ t := 0;\ (ctrl_o;\ dyn)^*)^* \text{ see Model 2} \tag{84}$$

Therefore, in Model 11 we now model the control of the obstacle in an inner loop around the continuous evolution *dyn* in (84) so that the obstacle control can interrupt continuous evolution at any time to make a decision, and then continue dynamics immediately without giving the robot's controller a chance to run. This means that the obstacle can make as many control decisions as it wants during the continuous evolution. The controller of the robot is still guaranteed to be invoked after at most time ε has passed, as modeled with the evolution domain constraint $t \leq \varepsilon$.

Theorem 10 (Passive safety for non-synchronized controllers). If started under initial conditions ϕ_{ps}, the model dw_{psns} (Model 11) guarantees passive safety, as expressed by the provable dL formula $\phi_{\text{ps}} \rightarrow [dw_{\text{psns}}]\psi_{\text{ps}}$.

We proved Theorem 10 with KeYmaera. The proof uses ϕ_{ps} as an invariant for the outer loop, whereas the invariant for the inner loop additionally preserves the differential invariants used for handling the dynamics *dyn*.

7.5 Arbitrarily Many Obstacles

The models so far are phrased in terms of a single but nondeterministic obstacle model, which is characterized by a maximum velocity V that an obstacle can have. This is safe, since all obstacles are considered by the nondeterministic obstacle choices. In the presence of different obstacles with different properties this can, however, lead to an inefficient controller, because we have to consider the same maximum velocity V for all obstacles, even if they differ significantly (e. g. stationary obstacles with maximal velocity $V = 0$ and pedestrians or even other cars with larger velocity bounds).

Modeling To make the controller more efficient, we will go from hybrid programs to quantified hybrid programs for distributed hybrid systems [28], i.e., systems that combine distributed systems aspects (lots of obstacles) with hybrid systems aspects (discrete control decisions and continuous motion). We introduce a sort O representing obstacles so that arbitrarily many obstacles can be represented in the model simultaneously. Each object of the sort O has the following properties: a maximum velocity, a position and an actual velocity. To model these properties, we introduce the non-rigid functions $p_o : O \rightarrow \mathbb{R}^2$, $v_o : O \rightarrow \mathbb{R}^2$, and $V : O \rightarrow \mathbb{R}$ such that $p_o(i)$ describes the current position of obstacle i and $v_o(i)$ the current velocity of

[14] Note that dL follows the common assumption that discrete actions do not take time; time only passes in ODEs. So all discrete actions happen at the same real point in time, even though they are ordered sequentially.

obstacle i as well as $V(i)$ the maximal velocity of a specific obstacle instance i. Note that $p_o(i)$ and $v_o(i)$ are vectors in two-dimensional space.

This new modeling paradigm also allows for another improvement in the model: So far, the closest obstacle was chosen by picking a position in ctrl_r nondeterministically. Such a nondeterministic choice includes the closest one. A real controller, however, needs to compute which obstacle is actually the closest one (or consider them all one at a time). Instead of choosing the closest obstacle nondeterministically in the model, using QdL we can consider all obstacles by quantifying over all instances of the sort O.

In the obstacle controller ctrl_o (86) we use a loop to allow multiple obstacles to make a control decision. Each run of that loop selects an obstacle instance i arbitrarily and updates its velocity vector (but no longer its position, since obstacles are now tracked individually). The loop can be repeated arbitrarily often. Thus, some arbitrary finite number of obstacles can make control choices in (86). In the continuous evolution, instead, we have to quantify over *all obstacles* in order to express that all obstacles change their state according to the differential equations simultaneously, see (89).

Model 12 Explicit representation of countably many obstacles, extends Model 2

$$\mathit{dw}_{\text{nobs}} \equiv (\mathit{ctrl}_o;\ \mathit{ctrl}_r;\ \mathit{dyn})^* \tag{85}$$

$$\mathit{ctrl}_o \equiv \big(i := *;\ v_o(i) := (*,*);\ ?\|v_o(i)\| \leq V(i)\big)^* \tag{86}$$

$$\mathit{ctrl}_r \equiv \text{see Model 2 on page 17 for details} \tag{87}$$

$$\mathit{safe} \equiv \forall i \in O.\ \|p_r - p_o(i)\|_\infty > \begin{cases} -\frac{v_r^2}{a_r} - V(i)\frac{v_r}{a_r} & \text{if } v_r + a_r\varepsilon < 0 \\ \frac{v_r^2}{2b} + V(i)\frac{v_r}{b} + \left(\frac{a_r}{b}+1\right)\left(\frac{a_r}{2}\varepsilon^2 + \varepsilon(v_r + V(i))\right) & \text{otherwise} \end{cases} \tag{88}$$

$$\mathit{dyn} \equiv \forall i \in O.\ (t' = 1,\ p_r' = v_r d_r,\ d_r' = -\omega_r d_r^\perp,\ v_r' = a_r,\ \omega_r' = \frac{a_r}{r},\ p_o'(i) = v_o(i)\ \&\ v_r \geq 0 \wedge t \leq \varepsilon) \tag{89}$$

We phrase the loop invariant φ_{nobs}, the initial condition ϕ_{nobs}, and the safety condition ψ_{nobs} as before except that they are now phrased for all obstacles $i \in O$, see (90)–(92). These generalize the single-obstacle conditions (30)–(32) on page 18.

$$\varphi_{\text{nobs}} \ \equiv\ v_r \neq 0 \rightarrow \forall i \in O. \|p_r - p_o(i)\|_\infty > \frac{v_r^2}{2b} + V(i)\frac{v_r}{b} \tag{90}$$

$$\phi_{\text{nobs}} \ \equiv\ v_r = 0 \wedge r_c \neq 0 \wedge \|d_r\| = 1 \tag{91}$$

$$\psi_{\text{nobs}} \ \equiv\ v_r \neq 0 \rightarrow \forall i \in O. \|p_r - p_o(i)\|_\infty > 0 \tag{92}$$

Verification These conditions are the basis for defining passive safety in the presence of arbitrary many obstacles in Theorem 11.

Theorem 11 (Passive safety in presence of arbitrary obstacles). If ϕ_{nobs} holds initially, then the model $\mathit{dw}_{\text{nobs}}$ (Model 12) guarantees passive safety, as expressed by the provable QdL formula $\phi_{\text{nobs}} \rightarrow [\mathit{dw}_{\text{nobs}}]\psi_{\text{nobs}}$.

We proved Theorem 11 with KeYmaera.

8 Liveness Verification of Ground Robot Navigation

Safety properties formalize that some bad behavior (such as collisions) will never happen. Liveness properties formalize that certain good things (such as reaching a goal) will ultimately happen. It would be easy to design a trivial controller that is only safe (just never move) or only live (full speed toward the goal ignoring all obstacles). The trick is to design robot controllers that meet both goals. The safe controllers identified in the previous sections are guaranteed safe (no collisions) and also allow motion. This combination of guaranteed safety under all circumstances (by a proof) and validated liveness under usual circumstances (validated only by some tests) is often sufficient for practical purposes. Yet, without a liveness proof, there is no guarantee that the robot controller will reach its respective goal except in the circumstances that have been tested before. In this section, we will, thus, verify liveness properties, since the precision gained by formalizing the desired liveness properties as well as the circumstances under which they can be guaranteed is quite insightful.

Formalizing liveness properties is even more difficult and the resulting questions are usually much harder than safety. Both safety and liveness properties only hold when they are true in the myriad of situations with different environmental behavior that they conjecture. They are diametrically opposed, because liveness requires motion but safety considerations may inhibit motion. For the safe robot models that we consider here, liveness is, thus, quite a challenge, because there are many ways that environmental conditions or behavior would make the robot stop or turn around for safety reasons, preventing it from achieving its goals. For example, an unrestricted obstacle could move around to just block the robot's path and then, as the robot re-plans to find another trajectory, move over to block the new path instead. To guarantee liveness, one has to characterize *all necessary conditions* that allow the robot to reach its goal, which are often prohibitively many. Full adversarial behavior can be handled but is challenging [32].

For a liveness proof, we deem three conditions important:

Adversarial behavior. Carefully defines acceptable adversarial behavior that the robot can handle. For example, crossing a robot's path might be acceptable in the operating conditions, but trapping the robot in a corner might not.

Conflicting goals. Identifies conflicting goals for different agents. For example, if the goal of one robot is to indefinitely occupy a certain space and that of another is to reach this very space there is no possible way for both to satisfy their requirements.

Progress. Characterizes progress formally. For example, in the presence of obstacles a robot sometimes needs to move away from the goal in order to ultimately get to the goal. But how far is a robot allowed to deviate on the detour?

Liveness conditions that are true, thus, usually need to define some reasonable restrictions on the behavior of adversarial agents in the environment. For example, a movable obstacle may block the robot's way for some limited amount of time, but not indefinitely. And when the obstacles moves on it may not turn around immediately again. Liveness conditions may also define a compromise between reaching the goal and having at least invested *reasonable effort* of trying to get to the goal, if unacceptable adversarial behavior occurred or goals conflicted, or progress is physically impossible.

In this section, we start with a completely stationary environment, so that we first can concentrate on finding a notion for progress. Next, we let obstacles cross the robot's path and define what degree of adversarial behavior is acceptable.

Model 13 Robot follows a straight line to reach a waypoint

$$\mathit{dw}_{\text{wp}} \equiv (\mathit{ctrl}; \mathit{dyn})^* \tag{94}$$

$$\mathit{ctrl} \equiv \quad (a_r := -b) \tag{95}$$

$$\cup\ (?v_r = 0;\ a_r := 0) \tag{96}$$

$$\cup\ (?p_r + \frac{v_r^2}{2b} + \left(\frac{A}{b} + 1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon v_r\right) < p_g + \Delta_g;\ a_r := *;\ ? - b \leq a_r \leq A); \tag{97}$$

$$\cup\ (?p_r < p_g - \Delta_g \wedge v_r \leq V_g;\ a_r := *;\ ? - b \leq a_r \leq \frac{V_g - v_r}{\varepsilon} \leq A) \tag{98}$$

$$\mathit{dyn} \equiv t := 0;\ \{p_r' = v_r,\ v_r' = a_r,\ t' = 1\ \&\ t \leq \varepsilon \wedge v_r \geq 0\} \tag{99}$$

8.1 Reach a Waypoint on a Straight Lane

As a first liveness property, we consider a completely stationary environment without obstacles, which prevents adversarial behavior as well as conflicting goals, so that we can concentrate on the conditions to describe how the robot makes progress without the environment interfering actively. We focus on the low-level motion planning where the robot has to make decisions about acceleration and braking in order to drive to a waypoint on a straight line. We want our robot to provably reach the waypoint, so that a high-level planning algorithm knows that the robot can reliably execute its plan by stitching together the complete path from straight-line segments between waypoints. We consider a simplified version where the robot has to stop at the waypoint, before it turns to head towards the next. That way, we can split a path into straight-line segments that make defining progress easier, because they are describable with solvable differential equations when abstracted into one-dimensional space.

Modeling We say that the robot reached the waypoint when it stops inside a region of size Δ_g around the waypoint. That is: (i) at least one execution enters the goal region, and (ii) all executions stop before exiting the goal region $p_g + \Delta_g$. The liveness property ψ_{wp} (93) characterizes these conditions formally.

$$\psi_{\text{wp}} \equiv \langle \mathit{dw}_{\text{wp}} \rangle (p_g - \Delta_g < p_r) \wedge [\mathit{dw}_{\text{wp}}](p_r < p_g + \Delta_g) \tag{93}$$

Remark 1. The liveness property ψ_{wp} (93) is formulated as a conjunction of two constraints: at least one run enters the goal region $\langle \mathit{dw}_{\text{wp}} \rangle p_g - \Delta_g < p_r$, while none exits the goal region on the other end $[\mathit{dw}_{\text{wp}}] p_r < p_g + \Delta_g$. In particular, there is a run that will stop inside the goal region, which, explicitly, corresponds to extending formula (93) to the following liveness property:

$$\begin{aligned} \langle \mathit{dw}_{\text{wp}} \rangle (p_g - \Delta_g < p_r \wedge 0 \leq v_r \leq V_g \wedge \langle \mathit{dw}_{\text{wp}} \rangle v_r = 0) \\ \wedge [\mathit{dw}_{\text{wp}}](p_r < p_g + \Delta_g) \end{aligned} \tag{100}$$

Formula (100) means that there is an execution of model dw_{wp} where the robot enters the goal region without exceeding the maximum approach velocity V_g, and from where the model has an execution that will stop the robot $\langle \mathit{dw}_{\text{wp}} \rangle v_r = 0$. The proof for formula (100) uses $v_r = 0 \vee (v_r > 0 \wedge v_r - n\varepsilon b \leq 0)$ to characterize progress (i. e., braking for duration $n\varepsilon$ will stop the robot).

Model 13 describes the behavior of the robot for approaching a goal region. In addition to the three familiar options from previous models of braking unconditionally (95), staying stopped (96), or accelerating when safe (97), the model now contains a fourth control option (98) to carefully approach the goal region.

Recall that in a proof, we have to show that the robot will get to the goal under any circumstance except those explicitly characterized as being assumed not to happen, e. g., unreasonably small goal regions, high robot velocity, or hardware faults, such as engine or brake failure. Similar to safety proofs, these assumptions are often linked. For example, a goal region is unreasonably small given a robot's engine parameters, if the robot will cross it without possibility to stop by just accelerating once from its initial position. In this case, both options of the robot will violate our liveness condition: it can either stay stopped and not reach the goal, or it can start driving and miss the goal.

Therefore, we introduce a maximum velocity V_g that the robot has to obey when it is close to the goal, similar to trains entering train stations slowly. That velocity must be small enough so that the robot can stop inside the goal region and is used as follows. While obeying the approach velocity V_g outside the goal region (98), the robot can choose any acceleration that will not let it exceed the maximum approach velocity. The dynamics of the robot in this model follows a straight line, assuming it is already oriented directly towards the goal (99).

Identification of Live Controls Now that we know what the goal of the robot will be, we give the intuition behind the conditions that make achieving the goal possible at all. The robot is only allowed to adapt its velocity with other choices than full braking when those choices are not going to overshoot the goal region, see $p_g + \Delta_g$ in (97) and $p_g - \Delta_g$ in (98). Condition (98) $-b \leq a_r \leq \frac{V_g - v_r}{\varepsilon} \leq A$ ensures that the robot will only pick acceleration values that will not exceed the approach velocity V_g in the next ε time units, i. e., until it can revise its decision. Once inside the goal region, the only choice remaining is to brake (since allowed unconditionally), which makes the robot stop close to the waypoint.

The robot is stopped initially ($v_r = 0$) someplace outside the goal region ($p_r < p_g - \Delta_g$), its brakes $b > 0$ and engine $A > 0$ are working,[15] and it has some known reaction time $\varepsilon > 0$, cf. (101).

$$\begin{aligned} \phi_{\text{wp}} \equiv v_r = 0 \land p_r < p_g - \Delta_g \land b > 0 \land A > 0 \\ \land \varepsilon > 0 \land 0 < V_g \land V_g \varepsilon + \frac{V_g^2}{2b} < 2\Delta_g \end{aligned} \tag{101}$$

Most importantly, the approach velocity V_g and the size of the goal region $2\Delta_g$ must be adequate. That way, we know that the robot has a choice to approach the goal with a velocity that fits to the size of the goal region.

Verification Similar to safety verification, for liveness verification we combine the initial condition ϕ_{wp} (101), the model dw_{wp} Model 13, and the liveness property ψ_{wp} (93) in a provable theorem Theorem 12.

Theorem 12 (Reach waypoint). If the robot starts in a state where ϕ_{wp} (101) holds, then there exists at least one execution of the control model dw_{wp} such that $p_g - \Delta_g < p_r$ holds, while for all executions $p_r < p_g + \Delta_g$ holds, as expressed by the provable dL formula $\phi_{\text{wp}} \rightarrow \psi_{\text{wp}}$, i. e., with ψ_{wp} from (93) expanded:

$$\phi_{\text{wp}} \rightarrow \big(\langle dw_{\text{wp}} \rangle (p_g - \Delta_g < p_r) \land [dw_{\text{wp}}](p_r < p_g + \Delta_g)\big) \ .$$

[15] For safety $A \geq 0$ is sufficient, but in order to reach a goal the robot must be able to accelerate to non-zero velocities.

We proved Theorem 12 using KeYmaera. In this proof, we need a variant that characterizes what it means to make progress towards reaching the goal region [23]. This induction principle for liveness is dual to a loop invariant for proving safety properties about loops by induction. If the progress measure indicates the goal would be reachable with n iterations of the main loop of Model 13, then we have to show that by executing the loop once we can get to a state where the progress measure indicates the goal would be reachable in the remaining $n-1$ loop iterations.

Informally, we can reach the goal if the robot currently has a positive velocity $0 < v_r$ and it either enters the goal region by just driving for time $n\varepsilon$ with that velocity or it is already at the goal, as summarized by the loop variant φ_{wp} (102).

$$\varphi_{\text{wp}} \equiv 0 < v_r \leq V_g \land p_g - \Delta_g < \begin{cases} p_r + n\varepsilon v_r & \text{if } p_r \leq p_g - \Delta_g \\ p_r & \text{otherwise} \end{cases} \tag{102}$$

Next, we analyze liveness in the presence of other moving agents.

8.2 Cross an Intersection

In this section, we prove liveness for scenarios in which the robot has to pass an intersection, while a moving obstacle may cross the robot's path, so that the robot may need to stop for safety reasons to let the obstacle pass. We want to prove that the robot has a way to successfully pass the intersection. The model captures the general case of a point-intersection with two entering roads and two exits at the opposing side, so that it subsumes any scenario where a robot and an obstacle drive straight to cross an intersection, as illustrated in Fig. 9.

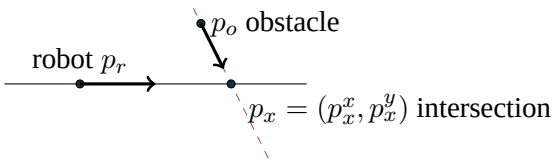

Figure 9: Illustration of the paths of a robot (black solid line) and an obstacle (red dashed line) crossing an intersection at point p_x.

Modeling Since there is a moving obstacle present, the robot needs to follow a collision avoidance protocol (here, for simplicity, we choose passive safety) in order to safely cross the intersection. Collision avoidance alone, however, will not reliably let the robot make progress. Thus, we will model a robot that favors making progress towards the other side of the intersection, and only falls back to collision avoidance when the obstacle is too close to pass safely.

Note, that intersections provide a simple strategy for the obstacle to prevent the robot from ever passing the intersection: the obstacle just needs to block the entire intersection forever by stopping there (e. g., somebody built a wall so that the intersection essentially disappeared). Clearly, no one could demand the robot passes the intersection in such impossible cases. We want to prove that the robot can pass the intersection when obstacles *behave reasonably*. We, therefore, have to include a restriction on how long the obstacle may reside at the intersection. For the sake of simplicity, we choose to introduce a strictly positive minimum velocity V_{min} to prevent the obstacle from stopping.

Model 14 Robot has to cross an intersection

$$dw_{\text{cx}} \equiv (ctrl_o;\ ctrl_r;\ dyn)^* \tag{103}$$

$$ctrl_o \equiv a_o := *;\ ?(-b \leq a_o \leq A) \tag{104}$$

$$ctrl_r \equiv \begin{cases} a_r := *;\ ? -b \leq a_r \leq A & \text{if } \textit{AfterX} \\ a_r := *;\ ?0 \leq a_r \leq A & \text{if } \textit{PassFaster} \\ a_r := 0 & \text{if } \textit{PassConst} \\ \quad (a_r := -b) & \\ \cup\ (?v_r = 0;\ a_r := 0) & \text{otherwise (cf. Model 2)} \\ \cup\ (?\textit{safe};\ \ldots) & \end{cases} \tag{105}$$

$$dyn \equiv t := 0;\ \{p_r' = v_r,\ v_r' = a_r, \tag{106}$$

$$p_o' = v_o,\ v_o' = a_o,\ t' = 1 \tag{107}$$

$$\&\ t \leq \varepsilon \wedge v_r \geq 0 \wedge v_o \geq V_{min}\} \tag{108}$$

Identification of Live Controls For ensuring progress, the model uses three conditions (*AfterX*, *PassFaster*, and *PassConst*) that tell the robot admissible conditions for choosing its acceleration, depending on its own position and the obstacle position in relation to the intersection. The robot can choose any acceleration if it already passed the intersection:

$$\textit{AfterX} \equiv p_r > p_x^x\ .$$

The robot is allowed to increase its speed if it manages to pass safely in front of the obstacle (even if the obstacle speeds up during the entire process), or if speeding up would still let the robot pass safely behind the obstacle (even if the obstacle drives with only minimum speed V_{min}):

$$\textit{PassFaster} \equiv v_r > 0 \wedge (\textit{PassFront} \vee \textit{PassBehind})$$

$$\textit{PassFront} \equiv p_o + v_o \frac{p_x^x - p_r}{v_r} + A\left(\frac{p_x^x - p_r}{v_r}\right)^2 < p_x^y$$

$$\textit{PassBehind} \equiv p_x^y < p_o + V_{min}\frac{p_x^x - p_r}{v_r + A\varepsilon}$$

The robot is allowed to just maintain its speed if it either passes safely in front or behind the obstacle with that speed:

$$\textit{PassConst} \equiv v_r > 0 \wedge p_x^y < p_o + V_{min}\frac{p_x^x - p_r}{v_r}\ .$$

In all other cases, the robot has to follow the collision avoidance protocol from Model 2 to choose its speed, modified accordingly for the one-dimensional variant here.

Verification As a liveness condition, we want to prove that the robot will make it past the intersection without colliding with the obstacle.

Theorem 13 (Pass Intersection). When starting with initial conditions ϕ_{cx} being satisfied, the robot can successfully pass the intersection $p_r > p_x^x$ when following dw_{cx} from Model 14:

$$\phi_{\text{cx}} \rightarrow \langle dw_{\text{cx}}\rangle p_r > p_x^x\ .$$

Model 15 Robot reaches a waypoint before a deadline expires

$$dw_{\text{wpdl}} \equiv (ctrl; dyn)^* \tag{110}$$

$$ctrl \equiv \begin{cases} (a_r := -b) \cup (?v_r = 0; a_r := 0) & \text{if } p_g - \Delta_g < p_r \\ a_r := A & \text{if } p_r + \frac{v_r^2 - V_g^2}{2b} + \left(\frac{A}{b} + 1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon v_r\right) \leq p_g - \Delta_g \\ a_r := *;\ ? - b \leq a_r \leq \frac{V_g - v_r}{\varepsilon} \leq A & \text{otherwise} \end{cases} \tag{111}$$

$$dyn \equiv t := 0;\ p_r' = v_r,\ v_r' = a_r,\ t' = 1,\ T' = -1\ \&\ t \leq \varepsilon \wedge v_r \geq 0 \tag{112}$$

We proved Theorem 13 in KeYmaera.

8.3 Liveness with Deadlines

So far, the liveness proofs showed that the robot is able to achieve a useful goal if it makes the right choices. The proofs do neither guarantee that the robot will always make the right decisions, nor do they specify how long it will take until the goal will be achieved. In this section, we want to ensure that it always achieves its goals within a reasonable amount of time. In other words: previously we showed that the robot can do the right thing to ultimately get to the goal, while here we want to ensure that it always makes the right decisions that will take it to the waypoint or let it cross an intersection within a bounded amount of time. This means, we need to show not only that there exists an execution such that the robot achieves its goals (ensures model feasibility), but also that all possible executions do so in the given time (ensures model correctness).

For such proofs, we need to specify a deadline. We are going to illustrate two alternatives for modeling deadlines: in Section 8.3.1 we will use a countdown T that is initialized to the deadline and expires when $T < 0$, whereas in Section 8.3.2 we will use T as a clock that is initialized to a negative starting value $T < 0$ and counts up to a deadline $D > 0$, so that two deadlines (crossing zero and exceeding D) can be represented with a single clock variable.

8.3.1 Reaching a Waypoint

Let us again start by defining a correctness condition for reaching a waypoint.

$$\psi_{\text{wpdl}} \equiv (T \leq 0 \rightarrow v_r = 0 \wedge p_g - \Delta_g < p_r) \wedge p_r < p_g + \Delta_g \tag{109}$$

Formula (109) expresses that after the deadline ($T \leq 0$, i. e. countdown T expired) the robot will be stopped inside the goal region ($v_r = 0 \wedge p_g - \Delta_g < p_r$). And that it will never be past the goal region ($p_r < p_g + \Delta_g$).

Modeling Model 15 is the familiar loop of control followed by dynamics (110). Unlike in previous models, braking and staying put is no longer allowed unconditionally for the sake of reaching the waypoint reliably in time (111). The robot accelerates maximally whenever possible without rushing past the waypoint region, cf. (111). In all other cases, the robot chooses acceleration to control towards the approach velocity V_g (111). The dynamics remain unchanged, except for the additional countdown $T' = -1$ of the deadline in (112).

Identification of Live Controls In order to prove this model live, we need to set achievable deadlines, as formalized in (113). The deadline has to be large enough (a) for the robot to accelerate to velocity V_g, (b) drive to the waypoint with that velocity, and (c) once it is there, still have sufficient time to stop. It also needs a slack time ε, so that the robot can react to being given the deadline. Finally, the conditions ϕ_{wp} that enable the robot to reach a waypoint at all have to hold as well.

$$\phi_{\text{wpdl}} \equiv T > \underbrace{\frac{V_g - v_r}{A}}_{(a)} + \underbrace{\frac{p_g - \Delta - p_r}{V_g}}_{(b)} + \underbrace{\frac{V_g}{b}}_{(c)} + \varepsilon \wedge \phi_{\text{wp}} \tag{113}$$

Verification A proof of the robot always making the right choices is a combination of a safety and a liveness proof: we have to prove that all choices of the robot reach the goal before the deadline expires (safety proof), and that there exists at least one way of the robot reaching the goal before the deadline expires (liveness proof).

Theorem 14. If the robot starts in a state where ϕ_{wpdl} (113) holds, then control model dw_{wpdl} ensures that the robot will reach the waypoint before the deadline expires, as expressed by the provable dL formula

$$\phi_{\text{wpdl}} \rightarrow \left([dw_{\text{wpdl}}]\psi_{\text{wpdl}} \wedge \langle dw_{\text{wpdl}}\rangle\psi_{\text{wpdl}}\right) \ .$$

We proved Theorem 14 using KeYmaera.

8.3.2 Crossing an Intersection

Crossing an intersection before a deadline is more complicated than reaching a waypoint, because the robot may need to wait for the intersection to clear so that it can cross safely.

Modeling Model 16 remains almost identical to Model 14, except for the robot controller, which has an additional control branch: when the obstacle already passed the intersection, we want the robot to pass as fast as it can by accelerating fully with maximum acceleration A.

Model 16 Crossing an intersection before a deadline expires

$$dw_{\text{cxd}} \equiv (ctrl_o;\ ctrl_r;\ dyn)^* \tag{114}$$

$$ctrl_o \equiv ctrl_o \text{ of Model 14} \tag{115}$$

$$ctrl_r \equiv \begin{cases} a_r := A & \text{if } p_o > p_x^y \\ ctrl_r \text{ of Model 14} & \text{otherwise} \end{cases} \tag{116}$$

$$dyn \equiv dyn \text{ of Model 14} \tag{117}$$

Identification of Live Controls Given the robot behavior of Model 16 above, we need to set a deadline that the robot can actually achieve, considering when (recall that it should still not collide with the obstacle) and how much progress the robot can actually make while driving. Hence, the deadline has to account for both the robot and the obstacle position relative to the intersection, as well as how much the robot can

accelerate. Let us start with the easiest case for finding a deadline D: when the obstacle already passed the intersection, the robot simply has to accelerate with maximum acceleration until it itself passes the intersection. The obstacles are assumed to never turn back, so accelerating fully is also a safe choice. The robot might be stopped. So, assuming we start a deadline timer T at time 0, the robot will drive a distance of $\frac{A}{2}D^2$ until the deadline D expires (i. e., until $T = D$). However, since we use a sampling interval of ε in the robot controller, the robot may not notice that the obstacle already passed the intersection for up to time ε, which means it will only accelerate for time $D - \varepsilon$. Formula (118) summarizes this case.

$$\eta^D_{\text{cxd}} \equiv D \geq \varepsilon \land p^x_x - p^x_r < \frac{A}{2}(D-\varepsilon)^2 \tag{118}$$

If unlucky, the robot determines that it cannot pass safely in front of the obstacle and will have to wait until the obstacle passed the intersection. Hence, in the deadline we have to account for the additional time that the obstacle may need at most to pass the intersection. We could increase D with the appropriate additional time and still start the timer at $T = 0$, if we rephrase the implicit definition of the deadline $p^x_x - p_r < \frac{A}{2}(D-\varepsilon)^2$ in (118) to its explicit form $D > \ldots$. In (119), instead, we start the deadline timer with a negative time $T < 0$,[16] such that it becomes $T = 0$ when the obstacle is located at the intersection.

$$\eta^T_{\text{cxd}} \;\equiv\; T = \min\left(0, \frac{p_o - p^y_x}{V_{min}}\right) \tag{119}$$

Verification Theorem 15 uses the deadline conditions (118) and (119) in a liveness property for Model 16.

Theorem 15 (Cross intersection before deadline). For appropriate deadline choices, there exists a run of model dw_{cxd}, such that when the deadline timer is expired ($T \geq D$) the robot is past the intersection ($p_r > p^x_x$).

$$\phi_{\text{cxd}} \land \eta^D_{\text{cxd}} \land \eta^T_{\text{cxd}} \to [dw_{\text{cxd}}](T \geq D \to p_r > p^x_x) \tag{120}$$

9 Interpretation of Verification Results

[16] Recall, that $p_o < p^y_x$ holds when the obstacle did not yet pass the intersection.

Table 5: Invariant and safety constraint summary

Safety	**Invariant** $+$ **Safe Control**
Static (Model 1, Theorem 1)	$\|p_r - p_o\|_\infty > \frac{v_r^2}{2b} + \left(\frac{A}{b}+1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon v_r\right)$
Passive (Model 2, Theorem 2)	$v_r \neq 0 \rightarrow \|p_r - p_o\|_\infty > \frac{v_r^2}{2b} + V\frac{v_r}{b} + \left(\frac{A}{b}+1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon(v_r+V)\right)$
Passive friendly (Model 3, Theorem 3)	$\|p_r - p_o\|_\infty > \frac{v_r^2}{2b} + V\left(\frac{v_r}{b}+\tau\right) + \frac{V^2}{2b_o} + \left(\frac{A}{b}+1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon(v_r+V)\right)$
Passive orientation	*isVisible* $> 0 \rightarrow \|p_r - p_o\|_\infty > \frac{v_r^2}{2b} + V\frac{v_r}{b} + \left(\frac{A}{b}+1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon(v_r+V)\right)$
(Model 5, Theorem 4)	and $\gamma\|r\| > \frac{v_r^2}{2b} + \left(\frac{A}{b}+1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon v_r\right)$
Extensions (passive safety examples)	
$+$ location uncertainty (Model 8)	$v_r \neq 0 \rightarrow \|\hat{p}_r - p_o\|_\infty > \frac{v_r^2}{2b} + V\frac{v_r}{b} + \left(\frac{A}{b}+1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon(v_r+V)\right) + \Delta_p$
$+$ actuator perturbation (Model 9)	$v_r \neq 0 \rightarrow \|p_r - p_o\|_\infty > \frac{v_r^2}{2b\Delta_a} + V\frac{v_r}{b\Delta_a} + \left(\frac{A}{b\Delta_a}+1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon(v_r+V)\right)$
$+$ velocity uncertainty (Model 10)	$v_r \neq 0 \rightarrow \|p_r - p_o\|_\infty > \frac{(\hat{v}+\Delta_v)^2}{2b} + V\frac{\hat{v}+\Delta_v}{b} + \left(\frac{A}{b}+1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon(\hat{v}+\Delta_v+V)\right)$
$+$ intermittent sensor failure	$v_r \neq 0 \rightarrow \|\hat{p}_r - p_o\| > \frac{v_r^2}{2b} + V\frac{v_r}{b} + \left(\frac{A}{b}+1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon(v_r+V)\right) + g\Delta_d$
with actual acceleration (Model 6)	$v_r \neq 0 \rightarrow \|p_r - p_o\|_\infty > \begin{cases} -\frac{v_r^2}{2a_r} - V\frac{v_r}{a_r} & \text{if } v_r + a_r\varepsilon < 0 \\ \frac{v_r^2}{2b} + V\frac{v_r}{b} + \left(\frac{a_r}{b}+1\right)\left(\frac{a_r}{2}\varepsilon^2 + \varepsilon(v_r+V)\right) & \text{otherwise} \end{cases}$

As part of the verification activity, we identified crucial safety constraints that have to be satisfied in order to choose a new curve or accelerate safely. These constraints are entirely symbolic and summarized in Table 5. When instantiated with concrete numerical values of a robot design, these safety constraints can be used to facilitate design decisions and get an intuition about how conservative or aggressive our robot can drive, such as

- how fast can the robot pass through a door?
- how fast can the robot drive on a corridor?

Below, we analyze the constraints for several values of acceleration force, braking force, control cycle time, and obstacle distance (i. e., door width, corridor width).

9.1 Safe Distances and Velocities

Static safety Let us recall the safe control constraint (13) from Model 1, which is justified by the proof of Theorem 1 to correctly capture when it is safe to accelerate in the presence of stationary obstacles p_o.

$$\|p_r - p_o\|_\infty > \frac{v_r^2}{2b} + \left(\frac{A}{b} + 1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon v_r\right)$$

The constraint links the current velocity v_r and the distance to the nearest obstacle through the design parameters A (maximum acceleration), b (maximum braking force), and ε (maximal controller cycle time). Table 6 lists concrete choices for these parameters and the minimum safety distance identified by (13) in Model 1. All except the third robot configuration (whose movement and acceleration capabilities outperform its reaction time) would lead to a reasonable performance in in-door navigation environments. Fig. 10 plots the minimum safety distance that a specific robot configuration requires in order to avoid stationary obstacles, obtained from (13) by instantiating the parameters A, b, ε and the current velocity v_r.

Table 6b turns the question around and lists concrete choices for these parameters and the resulting maximum safe velocity of the robot that (13) identifies. Fig. 11 plots the maximum velocity that the robot can travel in order to avoid stationary obstacles. The maximum velocity is obtained from (13) by instantiating

Table 6: Static safety: minimum safe distance and maximum velocity for select configurations

(a) Minimum safe distance

$v_r\ [\frac{m}{s}]$	$A\ [\frac{m}{s^2}]$	$b\ [\frac{m}{s^2}]$	$\varepsilon\ [s]$	$\|\mathbf{p_r} - \mathbf{p_o}\|\ [m]$
1	1	1	0.05	**0.61**
0.5	0.5	0.5	0.025	**0.28**
2	2	2	0.1	**1.42**
1	1	2	0.05	**0.33**
1	2	1	0.05	**0.66**

(b) Maximum velocity through corridors and doors

	$A\ [\frac{m}{s^2}]$	$b\ [\frac{m}{s^2}]$	$\varepsilon\ [s]$	$\mathbf{v_r}\ [\frac{m}{s}]$
Corridor $\|p_r - p_o\| = 1.25m$	1	1	0.05	**1.48**
	0.5	0.5	0.025	**1.09**
	2	2	0.1	**1.85**
	1	2	0.05	**2.08**
	2	1	0.05	**1.43**
Door $\|p_r - p_o\| = 0.25m$	1	1	0.05	**0.61**
	0.5	0.5	0.025	**0.47**
	2	2	0.1	**0.63**
	1	2	0.05	**0.85**
	2	1	0.05	**0.56**

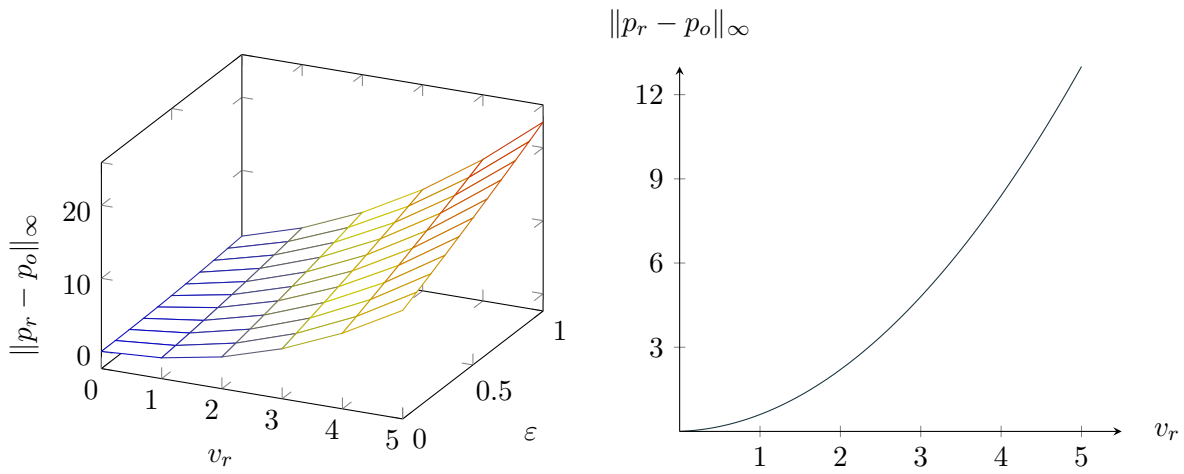

(a) By velocity and control cycle

(b) By velocity

Figure 10: Safety distance for static safety

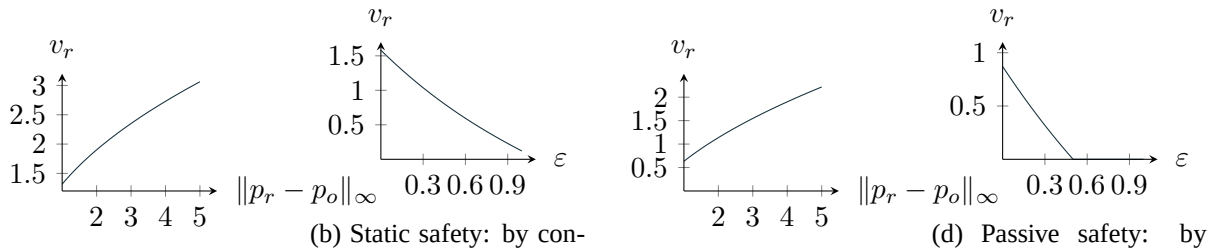

(a) Static safety: by distance for $\varepsilon = 0.05$

(b) Static safety: by control cycle time for $\|p_r - p_o\|_\infty = 1$

(c) Passive safety: by distance for $\varepsilon = 0.05$

(d) Passive safety: by control cycle time for $\|p_r - p_o\|_\infty = 1$

Figure 11: Comparison of safe velocities for static safety and passive safety with acceleration $A = 1$ and braking $b = 1$

the parameters A, b, ε and the distance to the nearest obstacle $\|p_r - p_o\|$. This way of reading the safety constraint (13) makes it possible to adapt the maximal desired velocity of the robot safely based on the current spatial relationships.

Moving obstacles Below, we repeat the safety constraint for accelerating or choosing a new curve in the presence of movable obstacles. The constraint introduces a new parameter V for the maximum velocity of obstacles.

$$\|p_r - p_o\|_\infty > \frac{v_r^2}{2b} + V\frac{v_r}{b} + \left(\frac{A}{b} + 1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon(v_r + V)\right)$$

Fig. 12 plots the minimum safety distance that the robot needs in order to maintain passive safety in the presence of moving obstacles. Note, that the maximum velocity in presence of movable obstacles can drop to zero when the obstacles can move too fast, the controller cycle time or the maximum acceleration force are too large, or when the maximum available braking force is too low.

Table 7: Passive safety: minimum safe distance and maximum velocity for select configurations

(a) Minimum safe distance

$v_r\ [\frac{m}{s}]$	$A\ [\frac{m}{s^2}]$	$b\ [\frac{m}{s^2}]$	$V\ [\frac{m}{s}]$	$\varepsilon\ [s]$	$\|\mathbf{p_r} - \mathbf{p_o}\|\ [m]$
1	1	1	1	0.05	**0.61**
0.5	0.5	0.5	0.5	0.025	**0.28**
2	2	2	2	0.1	**1.42**
1	1	2	1	0.05	**0.33**
1	2	1	2	0.05	**0.66**

(b) Maximum velocity through corridors and doors

	$A\ [\frac{m}{s^2}]$	$b\ [\frac{m}{s^2}]$	$V\ [\frac{m}{s}]$	$\varepsilon\ [s]$	$\mathbf{v_r}\ [\frac{m}{s}]$
Corridor $\|p_r - p_o\| = 1.25m$	1	1	1	0.05	**0.77**
	0.5	0.5	0.5	0.025	**0.69**
	2	2	2	0.1	**0.61**
	1	2	1	0.05	**0.4**
	2	1	2	0.05	**1.3**
Door $\|p_r - p_o\| = 0.25m$	1	1	1	0.05	**0.12**
	0.5	0.5	0.5	0.025	**0.18**
	2	2	2	0.1	**0**
	1	2	1	0.05	**0.26**
	2	1	2	0.05	**1**

9.2 Robot and Obstacle Shape

The safety analyses in Section 6 model the robots and the obstacles as points. This abstraction makes safety verification easier by reducing the number of variables. Safety results for different shapes of robots and obstacles can be obtained by the usual transformations [18] that account for the robot's shape as part of the obstacle (i. e., the robot's shape blows up every point of an obstacle). That way, distance measurements carried out by the robot at runtime are translated such that the measurements represent worst-case points. This method [18] transforms every obstacle point into a region that depends on the robot's shape; this is especially useful for distance measurements that provide point clouds as environment measurements, such as laser range finders. A trajectory of the robot is safe if it does not intersect any of the transformed regions.

10 Monitoring for Compliance At Runtime

Section 6–Section 8 discussed models of obstacle avoidance control and of the physical behavior of ground robots in their environment, and we proved that these models are guaranteed to possess safety and liveness properties. The proofs in Section 6–8 present strong evidence of the correctness of the models. If the models used for verification are an adequate representation of the real robot and its environment, these proofs transfer to the real system. But any model we can possibly build necessarily deviates from the real system at least to some extent.

In this section, we discuss how to use ModelPlex [20] to detect and safely respond to deviations between the model and the real robot in its environment at runtime. ModelPlex complements offline proofs with runtime monitoring: it periodically executes a *monitor*, which is systematically synthesized from the verified models by an automatic proof, and checks input from sensors and output to actuators for compliance with the model. If a deviation is detected, ModelPlex initiates a fail-safe action, such as stopping the robot to avoid actively running into obstacles, and, by that, ensure that *safety proofs* from the model carry over to the real robot.

A monitor checks actual evolution to discover failures and mismatches with the model. The acceleration chosen by the robot must fit to the current situation (for example, accelerate only when safe), the chosen curve must fit to the current orientation, and no unintended change to the robot's speed, position, orientation,

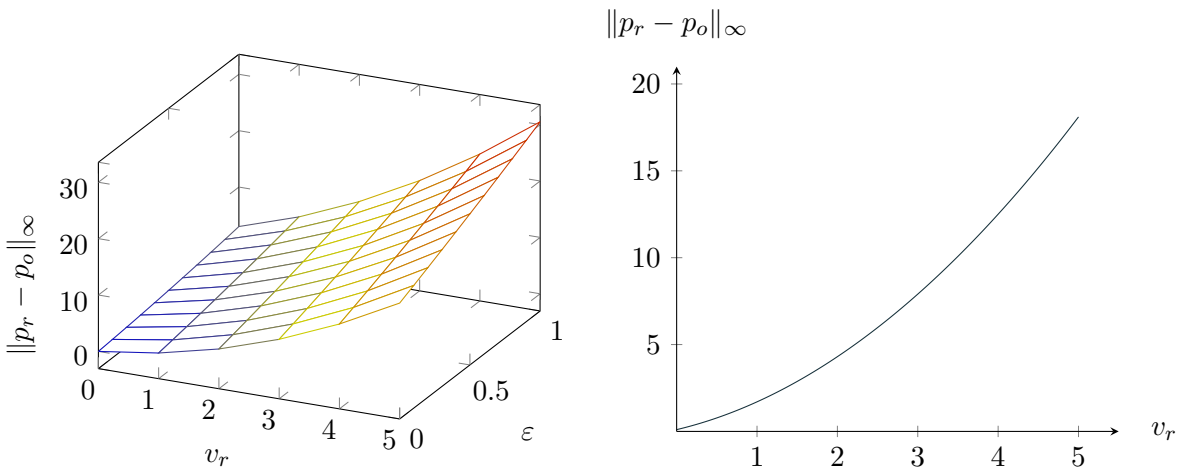

(a) By velocity and control cycle (b) By velocity

Figure 12: Safety distance for passive safety

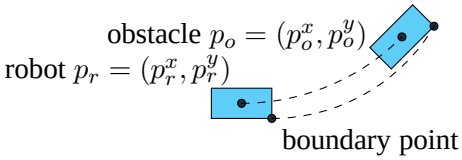

Figure 13: Each obstacle point becomes enlarged by the robot shape, transformed by a rotation and translation back to the robot shape [18]

or assumptions about the obstacles occurred. This means, any variable that is allowed to change in the model must be monitored. In the examples here, these variables include the robot's position p_r, longitudinal speed v_r, rotational speed ω_r, acceleration a_r, orientation d_r, curve r_c, and obstacle position p_o.

A ModelPlex monitor is designed for periodic sampling: for each variable there will be two observed values, one from the previous sample time (for example, robot position p_r) and one from the current sample time (for example, robot position p_r^+). It is not important for ModelPlex that the values are apart by exactly the sampling period, but merely that there is an upper bound (ε). A ModelPlex monitor checks in a provably correct way whether the evolution observed in the difference of the sampled values can be explained by the model. Fig. 14 illustrates the principle behind a ModelPlex monitor: the values from the previous sample time serve as starting state for executing the model. The values produced by executing the model are then compared to the values observed in the current sample time.

The verified models themselves, however, are not helpful as fast executable models, because they involve nondeterminism and differential equations. Hence, provably correct monitor expressions in real arithmetic are synthesized from a model using an offline proof in KeYmaera X. These expressions exhaustively capture the behavior of the models, projected onto the pairwise comparisons of sampled values that are needed at runtime. The full process is described in detail in [20].

Here, let us focus on a controller monitor expression synthesized from Model 2 above, which captures all possible decisions of the robot that are considered safe. A controller monitor [20] checks the decisions of an (unverified) controller implementation for being consistent with the discrete model *ctrl*. ModelPlex automatically obtains the discrete model from Model 2 with the ordinary differential equation (ODE) being

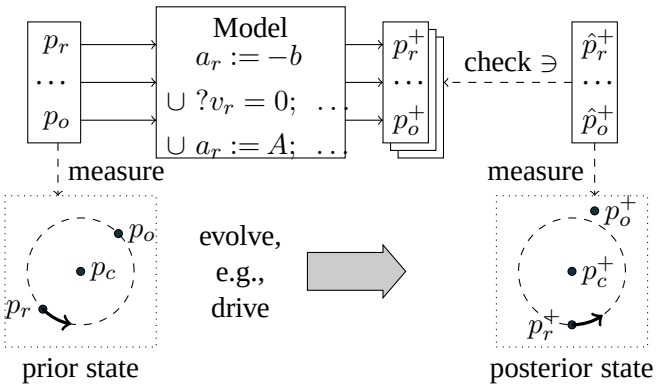

Figure 14: The principle behind a ModelPlex monitor: can the model reproduce or explain the observed real-world behavior?

safely over-approximated by its evolution domain. The resulting condition *monitor* in Fig. 15 (121), which is synthesized by a proof, follows the structure of the model: it captures the assumptions on the obstacle mon_o, the evolution domain from dynamics mon_{dyn}, as well as the specification for each of the three controller branches (braking mon_b, staying stopped mon_s, or accelerating mon_a).

The obstacle monitor part mon_o, see (122), says that the measured obstacle velocity d_r^+ must not exceed the assumptions made in the model about the maximum velocity of obstacles. The dynamics monitor part mon_{dyn}, see (123), checks the evolution domain of the ODE and that the controller did reset its clock ($t^+ = 0$). The braking monitor mon_b, see (124) defines that in emergency braking the controller must only hit the brakes and not change anything else (acceleration $a_r^+ = -b$, while everything else is of the form $x^+ = x$ meaning that no change is expected).[17] When staying stopped mon_s, see (125), the current robot speed must be zero ($v_r = 0$), and the controller must choose no acceleration and no rotation ($a_r = 0$ and $\omega_r = 0$), while everything else is unchanged. Finally, the acceleration monitor mon_a, see (126)–(127), when the distance is safe the robot can choose any acceleration in the physical limits $-b \leq a_r^+ \leq A$, a new non-spinning steering $c_r^+ \neq 0$ that fits to the current speed $\omega_r^+ c_r^+ = v_r$; position, orientation, and speed must not be set by the controller (those follow from the acceleration and steering choice).

11 Conclusion and Future Work

Robots are modeled by hybrid systems, because they share continuous physical motion with advanced computer algorithms controlling their behavior. We demonstrate that this understanding also helps proving robots safe. We develop hybrid system models of the dynamic window algorithm for autonomous ground vehicles and prove that the algorithm guarantees both passive safety and passive friendly safety in the presence of moving obstacles. All the proofs were achieved automatically with some manual guidance using the theorem prover KeYmaera. We interactively provided inductive and differential invariants (crucial) in KeYmaera, as well as reduced the number of terms and variables with simple interactions to reduce arithmetic complexity. 85% of the proof steps were automatic. Most interactive steps were simple arithmetic simplifications that KeYmaera's current proof strategies do not yet automate. We augment these models

[17] Note that unchanged obstacle position $p_r^+ = p_r$ means that the robot should not waste time measuring the obstacle's position, since braking is safe in any case.

$$\mathit{monitor} \equiv \mathit{mon}_o \wedge \mathit{mon}_{\mathit{dyn}} \wedge (\mathit{mon}_b \vee \mathit{mon}_s \vee \mathit{mon}_a) \tag{121}$$

$$\mathit{mon}_o \equiv \|d_r^+\| \leq V \tag{122}$$

$$\mathit{mon}_{\mathit{dyn}} \equiv 0 \leq \varepsilon \wedge v_r \geq 0 \wedge t^+ = 0 \tag{123}$$

$$\mathit{mon}_b \equiv p_o^+ = p_o \wedge p_r^+ = p_r \wedge d_r^+ = d_r \wedge v_r^+ = v_r \wedge \omega_r^+ = \omega_r \wedge a_r^+ = -b \wedge c_r^+ = c_r \tag{124}$$

$$\mathit{mon}_s \equiv v_r = 0 \wedge p_o^+ = p_o \wedge p_r^+ = p_r \wedge d_r^+ = d_r \wedge v_r^+ = v_r \wedge \omega_r^+ = 0 \wedge a_r^+ = 0 \wedge c_r^+ = c_r \tag{125}$$

$$\mathit{mon}_a \equiv -b \leq a_r^+ \leq A \wedge c_r^+ \neq 0 \wedge \omega_r^+ c_r^+ = v_r \wedge p_r^+ = p_r \wedge d_r^+ = d_r \wedge v_r^+ = v_r \tag{126}$$

$$\wedge \|p_r - p_o^+\|_\infty > \frac{v_r^2}{2b} + V\frac{v_r}{b} + \left(\frac{A}{b} + 1\right)\left(\frac{A}{2}\varepsilon^2 + \varepsilon(v_r + V)\right) \tag{127}$$

Figure 15: Synthesized safety conditions. The generated monitor captures conditions on obstacles mon_o, on dynamics $\mathit{mon}_{\mathit{dyn}}$, and on the robot's decisions on braking mon_b, staying stopped mon_s, and accelerating mon_a. The monitor distinguishes two observed values per variable, separated by a controller run (for example, p_r denotes the position before running the controller, whereas p_r^+ denotes the position after running the controller).

and safety proofs with robustness for localization uncertainty and imperfect actuation. We observe that incremental revision of models and proofs helps reducing the verification complexity and understanding the impact of uncertainty on the behavior of the robot. All parameters in our models—such as those for maximum obstacle velocity and sensor/actuator uncertainty—are fully symbolic and can be instantiated arbitrarily, including bounds from probabilistic models (e. g., assume the 2σ confidence interval of the distribution of obstacle velocities as maximum obstacle velocity). In this case, our verified safety guarantees translate into a safety probability.

Future work includes exploiting more kinematic capabilities (e. g., going sideways with omni-drive) and explicit accounts for distance measurement uncertainty, which is, however, easier than location uncertainty.

References

[1] Daniel Althoff, James J. Kuffner, Dirk Wollherr, and Martin Buss. 2012. Safety assessment of robot trajectories for navigation in uncertain and dynamic environments. *Autonomous Robots* 32 (2012), 285–302. DOI:http://dx.doi.org/10.1007/s10514-011-9257-9

[2] Francisco Bonin-Font, Alberto Ortiz, and Gabriel Oliver. 2008. Visual Navigation for Mobile Robots: A Survey. *J. Intelligent Robotics Systems* 53, 3 (2008), 263–296. DOI: http://dx.doi.org/10.1007/s10846-008-9235-4

[3] Sara Bouraine, Thierry Fraichard, and Hassen Salhi. 2012. Provably safe navigation for mobile robots with limited field-of-views in dynamic environments. *Autonomous Robots* 32, 3 (2012), 267–283. DOI:http://dx.doi.org/10.1007/s10514-011-9258-8

[4] Thomas Bräunl. 2006. Driving Robots. In *Embedded Robotics: Mobile Robot Design and Applications with Embedded Systems*. Springer, 97–111.

[5] Oliver Brock and Oussama Khatib. 1999. High-speed navigation using the global dynamic window approach. In *Robotics and Automation*. IEEE, 341–346. DOI: http://dx.doi.org/10.1109/ROBOT.1999.770002

[6] Howie Choset, Kevin Lynch, Seth Hutchinson, George Kantor, Wolfram Burgard, Lydia Kavraki, and Sebastian Thrun. 2005. *Principles Of Robot Motion*. MIT Press.

[7] Paolo Fiorini and Erwin Prassler. 2000. Cleaning and Household Robots: A Technology Survey. *Autonomous Robots* 9 (2000), 227–235. DOI: http://dx.doi.org/10.1023/A:1008954632763

[8] Paolo Fiorini and Zvi Shiller. 1998. Motion Planning in Dynamic Environments Using Velocity Obstacles. *J. Robotics Research* 17, 7 (1998), 760–772. DOI: http://dx.doi.org/10.1177/027836499801700706

[9] Dieter Fox, Wolfram Burgard, and Sebastian Thrun. 1997. The dynamic window approach to collision avoidance. *Robotics Automation Magazine, IEEE* 4, 1 (1997), 23–33. DOI: http://dx.doi.org/10.1109/100.580977

[10] Goran Frehse, Colas Le Guernic, Alexandre Donzé, Scott Cotton, Rajarshi Ray, Olivier Lebeltel, Rodolfo Ripado, Antoine Girard, Thao Dang, and Oded Maler. 2011. SpaceEx: Scalable Verification of Hybrid Systems. In *CAV (LNCS)*, S. Qadeer G. Gopalakrishnan (Ed.). Springer.

[11] Peter E. Hart, Nils J. Nilsson, and Bertram Raphael. 1968. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Trans. Systems Science and Cybernetics* 4, 2 (1968), 100–107. DOI:http://dx.doi.org/10.1109/TSSC.1968.300136

[12] Oussama Khatib. 1985. Real-time obstacle avoidance for manipulators and mobile robots. In *Robotics and Automation*. 500–505. DOI:http://dx.doi.org/10.1109/ROBOT.1985.1087247

[13] Steven M. LaValle and James J. Kuffner Jr. 2001. Randomized Kinodynamic Planning. *I. J. Robotic Res.* 20, 5 (2001), 378–400. DOI:http://dx.doi.org/10.1177/02783640122067453

[14] Sarah M. Loos, André Platzer, and Ligia Nistor. 2011. Adaptive Cruise Control: Hybrid, Distributed, and Now Formally Verified. In *FM*. Springer, 42–56. DOI: http://dx.doi.org/10.1007/978-3-642-21437-0_6

[15] Sarah M. Loos, David Renshaw, and André Platzer. 2013a. Formal Verification of Distributed Aircraft Controllers. In *HSCC*. ACM.

[16] Sarah M. Loos, David Witmer, Peter Steenkiste, and André Platzer. 2013b. Efficiency Analysis of Formally Verified Adaptive Cruise Controllers. In *ITSC*, Andreas Hegyi and Bart De Schutter (Eds.). 1565–1570. DOI:http://dx.doi.org/10.1109/ITSC.2013.6728453

[17] Kristijan Maček, Dizan Alejandro Vasquez Govea, Thierry Fraichard, and Roland Siegwart. 2009. Towards Safe Vehicle Navigation in Dynamic Urban Scenarios. *Automatika* 50, 3–4 (2009), 184–194.

[18] Javier Minguez, Luis Montano, and José Santos-Victor. 2006. Abstracting Vehicle Shape and Kinematic Constraints from Obstacle Avoidance Methods. *Autonomous Robots* 20, 1 (2006), 43–59. DOI: http://dx.doi.org/10.1007/s10514-006-5363-5

[19] Stefan Mitsch, Sarah M. Loos, and André Platzer. 2012. Towards formal verification of freeway traffic control. In *ICCPS*. IEEE, 171–180. DOI:http://dx.doi.org/10.1109/ICCPS.2012.25

[20] Stefan Mitsch and André Platzer. 2016. ModelPlex: Verified Runtime Validation of Verified Cyber-Physical System Models. *Form. Methods Syst. Des.* (2016). DOI: http://dx.doi.org/10.1007/s10703-016-0241-z Special issue of selected papers from RV'14.

[21] Andreas Müller, Stefan Mitsch, Werner Retschitzegger, Wieland Schwinger, and André Platzer. 2016. A Component-based Approach to Hybrid Systems Safety Verification. In *IFM (LNCS)*, Erika Abraham and Marieke Huisman (Eds.). Springer.

[22] Jia Pan, L. Zhang, and D. Manocha. 2012. Collision-free and smooth trajectory computation in cluttered environments. *J. Robotics Research* 31, 10 (2012), 1155–1175. DOI: http://dx.doi.org/10.1177/0278364912453186

[23] André Platzer. 2008. Differential Dynamic Logic for Hybrid Systems. *J. Autom. Reas.* 41, 2 (2008), 143–189. DOI:http://dx.doi.org/10.1007/s10817-008-9103-8

[24] André Platzer. 2010a. Differential-algebraic Dynamic Logic for Differential-algebraic Programs. *J. Log. Comput.* 20, 1 (2010), 309–352. DOI: http://dx.doi.org/10.1093/logcom/exn070

[25] André Platzer. 2010b. *Logical Analysis of Hybrid Systems: Proving Theorems for Complex Dynamics*. Springer.

[26] André Platzer. 2010c. Quantified Differential Dynamic Logic for Distributed Hybrid Systems. In *CSL (LNCS)*, Anuj Dawar and Helmut Veith (Eds.), Vol. 6247. Springer, 469–483. DOI: http://dx.doi.org/10.1007/978-3-642-15205-4_36

[27] André Platzer. 2011. Stochastic Differential Dynamic Logic for Stochastic Hybrid Programs. In *CADE*. 431–445. DOI:http://dx.doi.org/10.1007/978-3-642-22438-6_34

[28] André Platzer. 2012a. A Complete Axiomatization of Quantified Differential Dynamic Logic for Distributed Hybrid Systems. *Logical Methods in Computer Science* 8, 4 (2012), 1–44. DOI: http://dx.doi.org/10.2168/LMCS-8(4:17)2012 Special issue for selected papers from CSL'10.

[29] André Platzer. 2012b. The Complete Proof Theory of Hybrid Systems. In *LICS*. IEEE, 541–550. DOI: http://dx.doi.org/10.1109/LICS.2012.64

[30] André Platzer. 2012c. Logics of Dynamical Systems. In *LICS*. IEEE, 13–24. DOI: http://dx.doi.org/10.1109/LICS.2012.13

[31] André Platzer. 2012d. The Structure of Differential Invariants and Differential Cut Elimination. *Logical Methods in Computer Science* 8, 4 (2012), 1–38. DOI: http://dx.doi.org/10.2168/LMCS-8(4:16)2012

[32] André Platzer. 2015. Differential Game Logic. *ACM Trans. Comput. Log.* 17, 1 (2015), 1:1–1:51. DOI:http://dx.doi.org/10.1145/2817824

[33] André Platzer and Edmund M. Clarke. 2009. Formal Verification of Curved Flight Collision Avoidance Maneuvers: A Case Study. In *FM*. Springer, 547–562. DOI: http://dx.doi.org/10.1007/978-3-642-05089-3_35

[34] André Platzer and Jan-David Quesel. 2008. KeYmaera: A Hybrid Theorem Prover for Hybrid Systems.. In *IJCAR (LNCS)*, Alessandro Armando, Peter Baumgartner, and Gilles Dowek (Eds.), Vol. 5195. Springer, 171–178.

[35] Shahar Sarid, Bingxin Xu, and Hadas Kress-Gazit. 2012. Guaranteeing High-Level Behaviors while Exploring Partially Known Maps. In *Robotics: Science and Systems*.

[36] Derek Seward, Conrad Pace, and Rahee Agate. 2007. Safe and effective navigation of autonomous robots in hazardous environments. *Autonomous Robots* 22 (2007), 223–242. DOI: http://dx.doi.org/10.1007/s10514-006-9721-0

[37] Holger Täubig, Udo Frese, Christoph Hertzberg, Christoph Lüth, Stefan Mohr, Elena Vorobev, and Dennis Walter. 2012. Guaranteeing functional safety: design for provability and computer-aided verification. *Autonomous Robots* 32 (2012), 303–331. DOI: http://dx.doi.org/10.1007/s10514-011-9271-y

[38] Jur van den Berg, Pieter Abbeel, and Ken Goldberg. 2010. LQG-MP: Optimized Path Planning for Robots with Motion Uncertainty and Imperfect State Information. In *Robotics: Science and Systems*.

[39] Albert Wu and Jonathan P. How. 2012. Guaranteed infinite horizon avoidance of unpredictable, dynamically constrained obstacles. *Autonomous Robots* 32 (2012), 227–242. DOI: http://dx.doi.org/10.1007/s10514-011-9266-8